

CLOUD MEMOIRS

VIEWS FROM BELOW,
INSIDE AND ABOVE

A look at the history of cloud computing
by the CohesiveFT team, industry
experts, and early adopters

cohesiveFT

By Margaret Walker, Ryan Koop, and Patrick Kerpan

Created October 2013
All Rights Reserved
CohesiveFT Copyright 2014

Cloud Memoirs: Views from Below, Inside, and Above

A look at the history of cloud computing by CohesiveFT team, industry experts, and early adopters.

Created October 2013

Edited by Margaret Walker, Ryan Koop, and Patrick Kerpan

CohesiveFT

All Rights Reserved. Copyright 2014.

Copyright © 2014 by CohesiveFT.

ISBN-13 978-0-615-95611-4

cohesiveFT

All Rights Reserved.

Published in the United States by Cohesive Flexible Technologies,
Corp.

www.cohesiveft.com

Clouds, Fog, and Flood - A Note On the Theme of the Book

CEO Patrick Kerpan



Clouds, fog and flood. These are the themes of this eBook. I used the phrase in one of my early industry predictions to forecast how both enterprise and consumer cloud adoption would lead the market into new realms. The phrase also represents how we think about our data, our products, our security, and our interactions within the cloud environment.

From the “cloud to the ground,” each chapter looks back at the short history of cloud computing from the early stages where clouds were far off, through the gathering clouds that turned into fog. Today that fog is shifting, and beginning to rain down the first drops of change.

We sifted through our blog posts from 2006 onward, without any editorial updates (although tempted), in order to present the posts as a reflection of ours and some of our colleagues’ thoughts at the time. Interspersed with each chapter are new, exclusive insights from friends, colleagues, and customers we met along the journey.

Clouds: Clouds mask the complex systems and processes that go into the features, functions, and access we seek from the likes of mobile banking platforms and in-depth business analytics. Cloud computing is alluring because it makes the computing resources and data centers needed at scale, in a place, or at a moment in time, on-demand, and instant-on. Rather than own and manage (and pay for utilities, rent, taxes and security guards) for a single organizations’ data center and hardware needs, companies can use the cloud model to scale computing needs to business needs.

Fog: Things got foggy as the industry grew and data, applications, and providers evolved. From Amazon East to hundreds of data centers in 23 countries from 42 providers, cloud technology matured. The creation, updating, uploading and syncing of cloud-based application data has created giant cloud data caches. This has ratcheted up the need for more and more cloud data storage and cloud computing capacity. Businesses, from online publishing, to investment tracking, to massively parallel game playing, now require the capabilities of cloud computing to access, store, and manage data on-demand. As the clouds and fog productively coalesced, the first drops of rain started to fall.

Flood: The flood is coming, and how we navigate the oncoming waters of new cloud-based products will reshape the direction of businesses, social interactions, and technology itself. Today, the cloud enables and drives a new flood of data and analysis that helps organizations make decisions, automate workflows, and build new business. This infusion of data, connectivity and computation into the fiber of all business means working with the cloud in all its forms has led us from “will businesses use the cloud” to “what business can you possibly do without the cloud?”

Cheers,

Pat K.

Cloud, Fog, and Flood: the Cloud Progression

Clouds on the Horizon: 2006 - 2008

Founders' Stories	10
Exclusive Viewpoint: The Beginning of CohesiveFT and the Cloud	11
Fred Hoch	
Welcome to the Elastic Server Beta	13
Patrick Kerpan - CohesiveFT Blog	
CohesiveFT's Elastic Server	15
Jeff Barr - AWS Blog	
All at once I saw a crowd	16
Patrick Kerpan - CohesiveFT Blog	
InfoWorld's top 10 tech startups for 2008	20
Bill Snyder - InfoWorld	
CloudCamp London: the inauguration	21
James Governor - RedMonk	
Win Win for Independent Software Vendors	23
Chris Purrington - CohesiveFT Blog	
Software Manufacturing: A CFT White Paper	25
Patrick Kerpan - CohesiveFT Blog	

Cloud Cover: 2008 - 2009

Exclusive Viewpoint: History and Transformation of the Cloud	28
Edmund J. Sutcliffe	
Exclusive Viewpoint: Using CohesiveFT and AWS to Prove Cloud Concepts	30
Chad Lawler	
We're #1, (and #2, kind of #3, and #6, and part of #10)	33
Patrick Kerpan - CohesiveFT Blog	
VPN Cubed – Cloud is Ready for the Enterprise	34
Krishnan Subramanian - Cloud Ave	

CohesiveFT VPN-Cubed: Not Your Daddy's Encrypted Tunnel	36
Chris Hoff - Rational Survivability	
2009 in Virtualization and Cloud Computing: The Year of the Virtualization Professional	39
Patrick Kerpan - VMBlog	
Fog Surrounds Cloud: 2009 - 2012	
Exclusive Viewpoint: Cloud Computing: Past, Present and Future	42
Krishnan Subramanian	
Exclusive Viewpoint: White Knight Solutions from CohesiveFT	45
James Elwood	
New VPN-Cubed Version and The Cloud Connectivity Market	46
Ryan Koop - CohesiveFT Blog	
2010 in Cloud Computing: GAME ON!	49
Patrick Kerpan - VMBlog	
Getting out of the weeds	51
Chris Swan - The State of Me	
Welcome to the User-Cloud (Part 1)	53
Patrick Kerpan - CohesiveFT Blog	
Welcome to the User-Cloud (Part 2)	56
Patrick Kerpan - CohesiveFT Blog	
AMQP – the enchanted corner of SOA	60
Chris Swan - The State of Me	
2011 in Cloud Computing - The Empire Strikes Back!	62
Patrick Kerpan - VMBlog	
The First Drops: Cloud in Production: 2012 - 2014	
Exclusive Viewpoint: Cloud Still Has a Long Way to Come	66
Chad Lawler	
After VMWorld, who's talking about the Application layer?	67
Margaret Walker - CohesiveFT Blog	

OpenFlow is Software Defined Networking, Software Defined Networking is not just OpenFlow	68
Patrick Kerpan - CohesiveFT Blog	
The future of your data: Cloud, Fog, and Flood	70
Patrick Kerpan - CohesiveFT Blog	
Security and Control in the Cloud: Three migration rules to break	73
Patrick Kerpan - Cloud Computing Journal	
Why you should Waste Infrastructure: DevTest is Dead	77
Margaret Walker - CohesiveFT Blog	
Software Defined Networking and Network Function Virtualization	79
Chris Swan - Wired Innovation Insights	
Cloud markets and networking layers	81
Patrick Kerpan - CohesiveFT Blog	
Our Market Landscape and an Awesome Prezi	83
Margaret Walker - CohesiveFT Blog	
The Future - Floods of Cloud Uses: 2014 onward	
2014 Predictions: Amazon will lead, but the gap is closing	86
The CohesiveFT Team - VMblog	
Exclusive Prediction: Cloud Computing in 2015	88
Patricia B. Seybold	
Exclusive Prediction: Little Cloud Quantum Houses	93
Patrick Kerpan	

Replication rights granted by Jeff Barr of AWS blog, James Governor of MonkChips, Krishnan Subraminian of Cloud Ave, Eric Knorr of InfoWorld, Chris Hoff of Rational Survivability, David Marshall of VMblog, Patrick Kerpan of PJKTECH blog, Chris Swan of The State of Me blog, Cloud Computing Journal, and Wired Innovation Insights.

Reproduction and republication subject to their approval.

Chapter photo credits: Wikimedia Commons

Copyright 2014, CohesiveFT

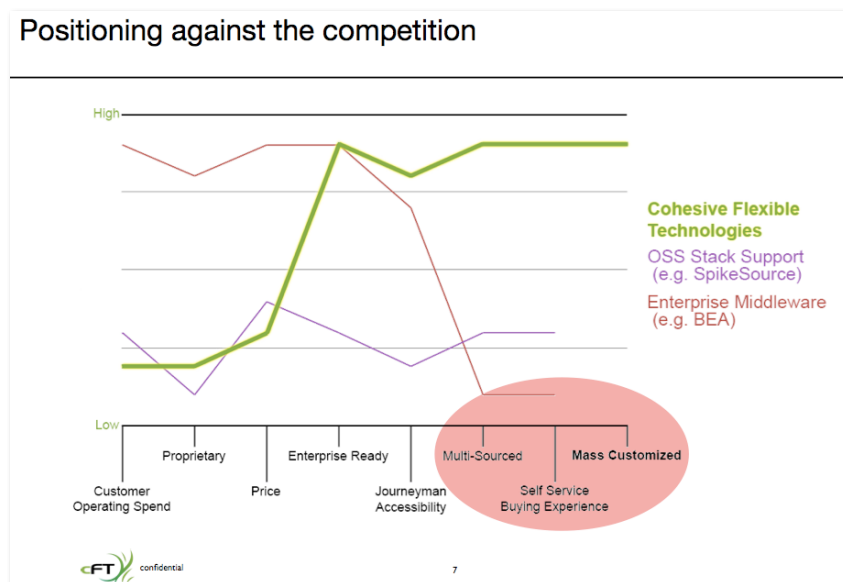
ISBN-13 978-0-615-95611-4



Clouds on the Horizon: 2006 - 2008

Founders' Stories

Patrick Kerpan, Craig Heimark, Dwight Koop, Alexis Richardson and Ryan Koop founded CohesiveFT after decades in enterprise IT and financial services management. CohesiveFT started in 2006 with venture backing and was originally focused on providing virtual appliance solutions for the financial services industry. The founders' backgrounds in networking, enterprise IT, and financial services technology allowed them to watch the industry "grow into cloud" from concept to reality.



The first network virtualization product CohesiveFT created was originally called vCubev, then VPN-Cubed and is now rebranded as VNS3. VNS3 is based on OpenVPN and Ubuntu, and was created as a solution to support the Server3 image management product. VNS3 connected the servers Server3 created together as if they were one logical group of resources. Customers began using VNS3 to connect networks to the cloud for internal and partner solutions in 2008. VNS3 allows applications to run unmodified as if they were all running on hosts behind a single switch. It works even when hosts are behind very restrictive firewalls, so it was a perfect fit for industries with regulation and data security concerns.

Since 2008, CohesiveFT have seen positive customer growth rate and increasing use across virtualized environments. As cloud infrastructure becomes more reasonably priced, more customers have sought out CohesiveFT through cloud providers' marketplaces and forums. VNS3 has helped users secure over 80 million virtual device hours in public, hybrid, and private clouds. From a handful of users in 2008, VNS3 now has over 700 customers in 25 countries.

Exclusive Viewpoint: The Beginning of CohesiveFT and the Cloud

Fred Hoch

President and CEO, Illinois Technology Association | Founder and Advisor, TechNexus



While the Cloud seems ubiquitous now, that wasn't always the case. In 2006, software as a service was still an early concern. I had just come to start the Illinois Technology Association (or ITA) after spending the last five years evangelizing the service movement. While the industry was certainly moving in that direction, customers were still hesitant. It was in this time frame that I met Pat, Dwight, and Craig...three financial services executives who had the vision of where things were going and how CohesiveFT was going to be a key part of that vision.

The software industry was ripe for change. The expense of enterprise IT has become too much for large corporations, and small corporations couldn't reap the benefits. Smart leaders saw that the movement to -as-a-Service was the answer. Cohesive saw that virtualization (now people call it "cloud") would become vital to businesses, and security would be the biggest challenge and opportunity.

I first saw what they were working on when they came by our ITA offices in early 2006. We were planning to use our space on Jefferson Street for events and promoting the tech industry in Illinois. But when the Cohesive team came by, we realized that we could let small companies share the space and build out their software right here. TechNexus grew from our first tenants to a much broader co-working and tech incubator space after that.

The most memorable thing from the early days of CohesiveFT at the original TechNexus location was in mid-2007. The team was working on a version of their first product, an application-specific virtual appliance for banks and high frequency trading. Then on May 24th 2007, Pat realized VMware was doing an application challenge to recognize and promote new apps for their Virtual Appliance Marketplace. They stayed in the offices almost all night to build and adapt their app to fit with VMware's marketplace. The product they build that night became "AppliaNCE" and did get certified later in August.

Only today is the market really catching up to what Cohesive saw back in 2006. Like most startups, the global recession taught everyone to be flexible. CohesiveFT later pivoted from a financial industry focus to making more multi-purpose tools for enterprises using cloud. Their first network software product, VcubeV, evolved into VPN-Cubed as they refocused, and now

they've adapted again and renamed it VNS3. I think this year will be the year that businesses understand the benefits and ease of using the cloud.

The Chicago startup scene, including TechNexus and the ITA, has been a great atmosphere to work in. Everyone from the mayor to local universities get involved in boosting local startups and promoting companies who succeed. We're on our way to becoming a rich technology community. The CohesiveFT team has really embraced our reputation of "honest Midwesterners" by creating real, valuable B2B software. It's been my pleasure to be part of that journey and I always find myself excited for the what's next.

Saturday, June 23, 2007

Welcome to the Elastic Server Beta

Patrick Kerpan - CohesiveFT Blog



CohesiveFT embarks today on an exciting extension to our existing business - the launch of our "Elastic Server On-Demand" service. Our new service which launches in Private Beta today, is the beginning of the realization of our vision of "making middleware accessible, manageable, and affordable".

As our company has grown we have noticed a key customer need - customers now want multi-sourced, loosely-coupled, vertically-aware middleware solutions. No longer is a "one size fits all" stack of computing capabilities acceptable. As a result, we are introducing our on-demand service with its ability to create "Elastic Servers™".

Elastic Servers are customer-configured, multi-sourced middleware stacks built dynamically in virtualization-ready formats and made available for download. Dynamic deployment capabilities to services like Amazon EC2 are coming quite soon.



Why do we call them "Elastic Servers"? Basically we are out to let customers build the stack they want - with no more - no less in it - and deployed in new dynamic forms like virtualization containers. Also, we are creating the ability for the server components to

flexibly link to each other, as well as other environmental services. We call these linkages "rubber bands". Over the next weeks as our beta rolls out, we will publish the formats needed to make components interact with services like our control panel or firewall.

We know we won't have a monopoly on interesting linkages, so end-user tools for creating components and rubber bands are high on our list. If you are already on the Private Beta you have seen the home directory of ESOD (elastic server on-demand) and its "micro-sites". These component-themed sites are just the beginning. End-user microsite tools are coming soon. In the mean time lets us know what components you would like made available, and what types of micro-sites are of interest to you.

What can these initial servers do? Right now we are beta testing the packaging, VM output formats, and user experience for ESOD. The middleware servers you create could have component conflicts, since for now we are building what you ask for - without checking a rule-

base of compatibility information. When you download and launch, your packages are there, but we haven't necessarily started their processes. To describe these first dynamically created servers as "Inert" middleware probably doesn't give us enough credit, "production ready" would give us too much. As our "rubber bands" come on line, and we get more of your feedback of what is needed, we will distance ourselves from the former description and move decisively towards the latter.

Take a look at the Elastic Server On-Demand roadmap and please take the feature poll to help us understand your needs. This is "your middleware". Let us know what you want to do with it and how we can help.

Oct 16 2007

CohesiveFT's Elastic Server

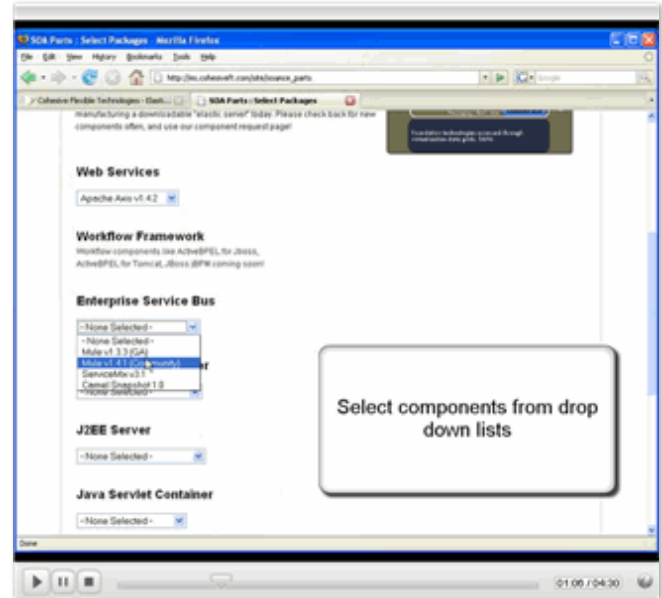
Jeff Barr - AWS Blog



As part of my most recent trip to London I paid a visit to a suite of offices shared by LShift and CohesiveFT. I gave them an in-depth AWS presentation and we also discussed Cohesive's

Elastic Server product.

Elastic Server simplifies the process of creating application stacks for use on EC2 and other virtual environments. Instead of the painstaking manual downloading, configuration, building, and installation that is otherwise required when creating a complex stack, Elastic Server provides a menu-driven selection and build process.



After choosing the desired web service environment, workflow framework, J2EE server, enterprise service bus, Java servlet container, page server, message queue and AMQP broker the user simply enters the desired virtualization parameters and initiates the build process. After a few minutes the build process spits out a finished EC2 AMI. From there it is easy to launch one or more instances of the AMI using the EC2 Firefox Manager or another EC2 tool. Each application stack is configured to include a copy of Cohesive's web-based Elastic Server Manager for easy management of each component.

If this sounds cool you should definitely watch the video to see how it works.

-- Jeff

The post originally appeared on the Amazon Web Services blog: <http://aws.typepad.com/aws/2007/10/cohesivefts-ela.html>

Monday, March 17, 2008

All at once I saw a crowd

Patrick Kerpan - CohesiveFT Blog



Last week James Governor of [Redmonk](#) published a blog entry entitled “15 Ways to Tell Its Not Cloud Computing”. It drew some very interesting commentary from Mark Cathcart, John M Willis, Chris Petrilli and Gerrit Huizenga.

James was kind enough to acknowledge my input in his blog entry and in return I’d like to offer some views on why I substantially agree with his blog post. Please note that this is an emerging space and the idea of a cloud remains nebulous. So what James is really saying, I think, is: let’s ask what we want from cloud computing in the next few years.

Clouds are all about delivering a “retail like” experience to users - even within the enterprise. Here is an analysis what this might mean.

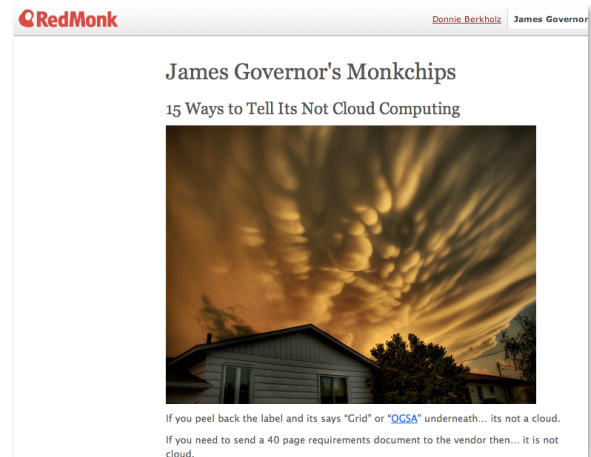
If you peel back the label and its says “Grid” or “OGSA” underneath... it’s not a cloud.

Clouds are about offering compelling economics for infrastructure as a service instead of as on premise servers or big desktops.

Clouds offer a fundamentally different abstraction than Grids. If there is a metaphorical ‘label’ then you should not be able to ‘peel it away’ at all. Grids provide a mechanism for running processes that implement a specific API, across multiple compute resources. Clouds provide a service for users to run whole applications by deploying them to a virtual data center.

Of course clouds can use a grid technology under the hood, for example to schedule slots. But the point is that many complexities are hidden from the user in order provide a clean and simple service. Whether the cloud is “public” like Amazon EC2, or in some sense an “enterprise” or “private” cloud, it delivers a “retail experience”. You should not need to read and understand multiple long documents to use it. I know of several hedge funds running grid style compute farms on EC2 using map/reduce on Xen images - it’s easier that way.

If you need to send a 40 page requirements document to the vendor then... it is not cloud.



In my opinion a cloud should offer “self service” capability. The cloud should not itself need to be customized by the provider, for individual customer use cases. This is not just about “productising” for the retail market. Consider an enterprise with two data centers. Clearly it saves money and time if applications can be migrated between data centers without refactoring. An enterprise knows it has implemented a cloud, when this is possible. Achieving this requires a certain abstraction.

If you can't buy it on your personal credit card... it is not a cloud.

Well - sort of. This is about managing cost when multiple users share common server infrastructure. Without a means to capture costs of use and attribute them to users, explaining where value comes from becomes a matter of guesswork. This hampers investment in new enterprise projects by staff who want access to shared infrastructure, and commonly leads to ‘guerilla’ use of things like Amazon EC2 by staff on their own credit cards.

The solution that I think the cloud model enforces is based on the discipline of a retail model. If data center providers made all users attribute costs as if they were paying by credit card, then everyone would know how much they were paying and for what. This would break down barriers to financing and starting new projects however small.

If they are trying to sell you hardware... it's not a cloud.

Clouds represent a way to separate ‘opex’ from ‘capex’. People making investment decisions about cloud resource needs should not need to factor in any amortized capital investments in hardware, because they only pay for what they need, when they need it. When a business case is being made for a new enterprise project, if hardware purchasing costs do not appear in the ROI calculation, then you probably have a cloud. Then, the data center managers decide how much hardware to buy and at what price to offer their cloud services to business units. An internal market may be used for dynamic pricing.

If there is no API... it's not a cloud.

If there is no API then it is not a service. Clouds virtualize resources (e.g. CPU, storage) as a service. Hence clouds must have APIs, e.g. a means to load and start up an application, or to stop it.

If you need to re-architect your systems for it... it's not a cloud.

The key word here is ‘need’. What we are discovering is that virtualizing whole compute environments may be the right way to ‘do’ utility computing. Whether computers are

delivered as bootable images or as VMs, this means that people are not required to completely rewrite their application to run it in the cloud. Do we need to explain why this is more attractive than the alternative? I didn't think so.

If it takes more than ten minutes to provision... it's not a cloud.

Is this controversial? As John M Willis points out, the whole notion of provisioning is being replaced by autonomics and elasticity for managing running systems. At CohesiveFT, we also argue that elasticity really means provisioning includes built-to-order system assembly.

If you can't de-provision in less than ten minutes... it's not a cloud.

This is really critical. Being able to stop something and clean it up, is much more important than being able to start it. Anything else just adds friction which translates into less incentive to use the service in the first place and hence less agility and other benefits.

If you know where the machines are... it's not a cloud.

Location matters for many applications such as trading systems, and complex clusters, because location impacts latency. But then the latency properties of a service ought to be exposed as part of its contract with the user. Then, perhaps users are willing to pay more for lower latency, or for better bit throughput.

What clouds ought to do, however, is get us away from any 'my rack, your rack' thinking because that breaks the retail model by recoupling opex to capex.

If there is a consultant in the room... it's not a cloud.

Of course if you want to build your own clouds and work out the business case for moving system administrators over to new roles such as "virtual sysadmin" then get a consultant.

For actual use of a retail-like service by the enterprise, or within the enterprise, you should not need a personal shopper unless your use of clouds is a vanity project.

If you need to specify the number of machines you want upfront... it's not a cloud.

The point is that you can increase your load whenever you like, so long as you can pay for it. Hence it does not matter if you start with zero machines or ten thousand. This is not the same as saying "you cannot request how many machines you need".

If it only runs one operating system... it's not a cloud.

This is about cloud providers not forcing end users to adopt their particular operating system distro and build. It characterizes the difference between the new 'self service' clouds based on virtualization technology, and the CPU-rental model of a couple of years ago.

Generally the people we meet have their own ideas about which O/S to use on their computers, so why not extend this to the cloud? The alternative is to offer a data center service which only runs a particular O/S build - this kind of Linux, that distribution, and that build. In an enterprise, this simply does not scale across multiple data centers over time. For a retail cloud, it's out of the question. This is why it is hard to see retail and enterprise clouds working commercially unless the compute unit is a virtual image or bootable image.

If you can't connect to it from your own machine... it's not a cloud.

Clouds are about consolidating resources in order to benefit from economies of scale. When a service is inaccessible or unavailable, then that diminishes its value to the service provider as well as the end user. We have found that people want to do things like spread their loads across their own machines, and data centers, or across multiple clouds securely. We call this 'multi-sourcing'. Obviously some degree of secure accessibility is a prerequisite for this.

If you need to install software to use it... it's not a cloud.

"You" means the user - clouds are about making life easier for users. If you are an enterprise with a cloud, or using clouds, then you may have teams all over the world, including folks coming and going all the time, or working from within IT suppliers or at customer sites, then your cloud will have a lot of users. You want to reduce any costs that you are currently paying for them to access your IT resources. This is why a 'retail experience' is so essential. Cloud means service.

If you own all the hardware... it's not a cloud.

Because you don't want to own all the hardware. It eats up power and space, gives off heat, and needs people to look after it. As more and more "green tape" appears in the form of environmental regulation, and if oil prices go higher, trust me you won't want to own and be responsible for owning and managing hardware.

If you need a three page blog post to explain it ... it's not a cloud

Enough already!

May 19, 2008

InfoWorld's top 10 tech startups for 2008

Bill Snyder - InfoWorld



Hot tech startup: Cohesive Flexible Technologies

Founded: 2006

Tech breakthrough: Software that builds a base image of a server and reformats it into a chosen virtualization configuration without first building a physical server.

Business problem addressed: How to manage and quickly deploy servers in a complex environment.

What the technology does: Patrick Kerpan, CTO of Cohesive Flexible Technologies, calls it the "more of everything problem." He notes, "Open source, open standards, virtualization, SOA and clouds are proliferating, needing countless components. An enterprise in the financial sector may have as many as 100 different application stacks." CohesiveFT's Elastic Server On Demand assembles virtual machines and deploys them to Amazon's Elastic Compute Cloud (EC2) or creates an implementation of major virtualization formats including EMC VMware, Citrix XenSource, and Parallels. CohesiveFT claims that a deployment can be completed in hours or even minutes. Using CohesiveFT's management system, IT can then track the deployed component assemblies throughout their lifecycle and log all configuration changes. The server can later be provisioned to another platform.

How the technology works: The company maintains libraries of components, including those from the open source community and software vendors. Customers may add their own proprietary components to the library (for their use only) and construct the image of a virtual application stack. The resulting images are built, encapsulated, given a unique identity and injected with management and integration services. CohesiveFT calls the completed stack an "elastic server." CohesiveFT uses an add-on to OpenVPN (open source virtual private network) that can connect multiple servers (both physical and virtual) located in various data centers and hosted at different providers into a single address space.

Forward spin: At the moment, Amazon's EC2 is the only cloud with a direct connection to CohesiveFT, but expect to see more clouds supported fairly soon, Kerpan says. The company also plans to add management tools as well as support for virtual Linux.

For the full list and the original article on InfoWorld, visit: <http://www.infoworld.com/t/hardware/infoworlds-top-10-tech-startups-2008-300?page=0,2>

July 21, 2008

CloudCamp London: the inauguration

James Governor - RedMonk



CloudCamp London rocked: around 250 people showed up, and the applause for the speakers was surprisingly generous from a geek UK audience; its clear there is a hunger for information in the space. The format wasn't perfect, but it was still a very good effort, even if there were no dead dragons lying bleeding afterwards. ...

The first speaker was Simon Wardley, otherwise known as the *AAS Master. Unlike some of the other speakers he actually stuck to the Lightning Talk format's 10 minute limit. Key phrase: "Yesterday's hot stuff becomes tomorrow's boredom". His presentation was the usual mashup of ducks and horror movie stills. The essential argument- unless open



source standards emerge, cloud will just be another round of industry lock in. One take on that argument is that customers always vote with their feet, and they tend vote for something somewhat proprietary – see Salesforce APEX and iPhone apps for example. Experience always comes before open. Even supposed open standards dorks these days are rushing headlong into the walled garden of gorgeousness we like call Apple Computer.

I think Will Fellows of the 451 Group did an excellent job of putting forward a cloud taxonomy. Really I was very impressed with the work he has put in. I would strongly advise you to ping him and ask for a free copy of the report his talk was based on (the offered one to cloudcamp attendees).

Last up was Alan Williamson. summarised eloquently here:

"His main message was that with cloud infrastructures problems don't magically go away, they just shift. You don't have scalability or storage problems any more, but you need to constantly monitor the cloud and your application in it. Alan pointed out examples when Amazon's cloud failed and their applications got cut off from the

Internet. As a solution, he proposed deploying the application on more than one cloud so that you have resilience. This requires writing the application in a way that can be easily ported to different providers, which in itself might be a challenge. One idea that was really striking was their analysis of getting off the cloud to a dedicated infrastructure again — apparently it would take them about three weeks of full-bandwidth transfer to download the data that they have in the cloud, making it virtually impossible to go back.”

Nice- so much for “freedom to leave“. The service might support it, but with massive data sets, portability ain’t so easy... Mi compadri Stephen O’Grady recently posted some good thoughts on Cloud Standards but its also worth considering the physical limits of data portability (we might be talking about flowing a terabyte of data, not just an email address). To often we assume everything on the web is instantaneous. We’re talking about the Physics, rather than the Economics of Data Portability. Data volumes will certainly be a key challenge for data portability, which is one reason my money is on the Synchronised Web.

Well perhaps not everything is instantaneous: as Will Fellows said, he was talking to a middleware vendor who said he wasn’t losing sales to Cloud Computing... but it was elongating the sales cycle. Not sure how that fits into 15 Ways To Tell Its Not Cloud Computing... I know there were quite a few suits at CloudCamp.

...

picture credit Chris Purrington. He says “all rights reserved” but I am sure he’ll consider this fair use.

The post was originally published by James Governor on Red Monk:<http://redmonk.com/jgovernor/category/cloud-computing/#ixzz2k5jZZ6mG>

Tuesday, August 5, 2008

Win Win for Independent Software Vendors

Chris Purrington - CohesiveFT Blog



Hi, I'm Chris Purrington. Not so long ago I was with Borland where I spent 5 years as UK Managing Director. Borland has a great heritage and a history of great products, you may know it in its current guise as the "Open ALM" company.

However, I saw it then and I see it now, regardless of whose tools are being used or how good they are, many customers are still frustrated that end-to-end the Software Development Life Cycle (SDLC) is still like an assault course, and very difficult for organisations to traverse. Communication barriers between roles throughout the lifecycle are still an issue today despite many ISVs looking to ease this pain.

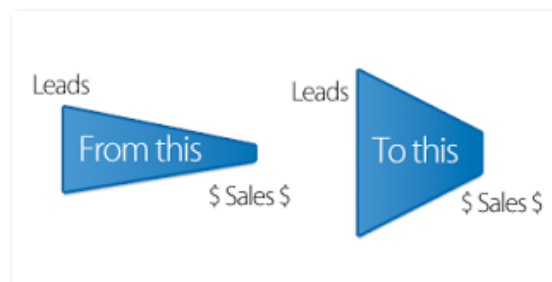
I can see CohesiveFT's Elastic Server changing the way organisations will tackle traditional challenges of SDLC. Fitting in alongside many other ALM tools working within existing processes, Elastic Server will deliver significant speed, quality, and cost benefits across all roles throughout the SDLC, acting as the glue in the hand-offs between functions.

For ISVs there's a Double Win

ISVs come in all shapes and sizes, and all can benefit from Elastic Server. I think the 'wins' are not just in Product Delivery and Data Centre Operations but also in the "sales cycle". ISVs can transform their whole pre-sales solution validation or Proof of Concept process, accelerating their sales success.

Wider, Shorter, Larger and more efficient!

That's the sort of sales pipeline I like. Elastic Server allows you to widen your pipeline funnel, shorten your sales cycle and close more deals. Your pipeline can go:



ELASTIC SEVER increases awareness, and makes solutions intuitive.

Ideally, enterprises customers would come to you after they know how to make best use of your product. But enterprise solutions have many capabilities, resulting in many use cases. Showing off all the relevant capabilities on offer to the appropriate people, can be a real challenge. Elastic Server makes it easy to create your own "component portals", each displaying representative use cases; use cases the customer can assemble and deploy at the click of a button.

Elastic Server allows you to easily create, save and make available (through your own micro portals) tailored instances of your solution. These are saved as templates and reused to create Elastic Sever versions of your solution whenever they are required. These Elastic Servers are custom application stacks, built from your solutions components, virtualization-ready, which means: ready to download or deploy to a cloud. So you can make as many instances as you need tailored or not, available on demand.

Your templates can be vanilla, or industry solution specific, or tailored to individual prospects. They can be easily reused, displayed, demoed and evaluated. Within your portals you can easily add your own branding, diagrams, and notes, effectively 'white-labeling' them, making your tailored micro portal more relevant to your prospect, more accessible and easier to follow.

Struggling with lengthy Proof Of Concepts?

From my Borland days I recall often just doing a POC required long procurement cycles and getting executive sign-off. If your POCs need several days or weeks of your prospect's architect, developer and data centre resources with a dedicated machine, then getting budget approval for a POC can be just as difficult as getting approval for a \$50k, or \$500k purchase. With Elastic Server and clouds such as Amazon EC2 or Flexiscale you can shrink this whole process to just one day, and no machine resources are required! So your sponsor doesn't need to go to the board or investment committee before they get a chance to prove the value of your solution. You can remove key barriers so your sales force can kick-off more POCs, faster.

A typical scenario of days or weeks to set up an enterprise POC, and having limited resource available can result in delays. In these situations it is easy for everything and everyone, to lose momentum. Imagine instead your pre-sales team selecting a template system that matches the prospects needs, then spending minutes personalising it for them, so within hours of agreeing to the POC you can email your prospect a link where they will find their trial system running.

And if your prospects want to run the POC in-house on their machines then you have the option to rebuild to any major virtualisation format from your Elastic Server templates.

Wednesday, September 3, 2008

Software Manufacturing: A CFT White Paper

Patrick Kerpan - CohesiveFT Blog



We just put the finishing touches on a white paper titled, "Software Manufacturing: Custom Application Stacks for Virtualized Infrastructure and Cloud Computing." (1.8MB .pdf). As you'll see, the cover is mostly white, and the document is indeed paper-shaped, as a white paper ought to be. For those

who are looking to minimize the number of words they take in, here's a quick high-level summary about the document.

Summary

Historically customers have acquired their middleware software on a vertically-integrated basis from one or a few category-dominant vendors. However, a customer revolution in service-oriented architecture (SOA) and service-oriented infrastructure (SOI), combined with ever more complex business needs, has driven customers to a multi-sourced middleware component architecture. As these market mechanics have evolved over the last decade, the middleware market is moving from single-sourced, tightly-coupled, vertically integrated middleware, to multi-sourced, loosely-coupled, vertically aware application and middleware stacks. While multi-sourcing allows customers to apply the latest middleware innovation to their business, it comes at the price of exploding complexity, as they manage many more combinations of available and applicable software components. Despite this, CohesiveFT (CFT) believes application stacks do not need to be complex to install, configure, deploy, or manage; whether to physical, virtual or cloud computing infrastructures.



CFT provides an on-demand platform that enables customers to build and manage application stacks for the virtualized data center; whether owned by the enterprise or "in the clouds." The Elastic Server™ platform provides the ability to assemble, manage and deploy a virtualized application stack that is dynamically defined from component libraries consisting of open source, commercial source, and proprietary customer code. It dramatically improves

the quality and consistency of application stack assemblies at the same time that it reduces time to market.

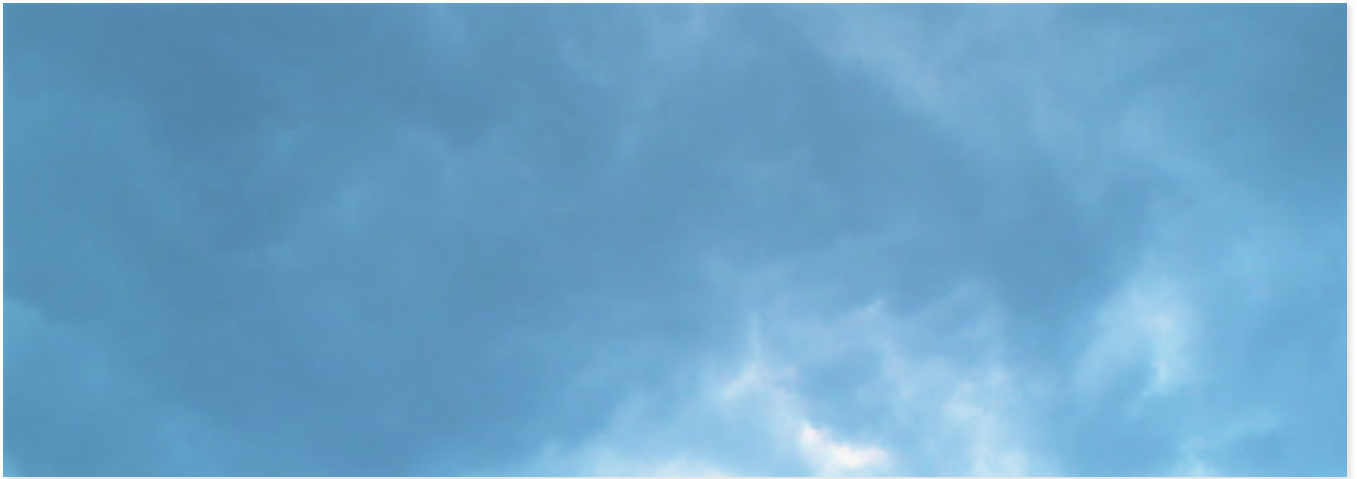
The Elastic Server platform allows customer choice of individual components at defined architectural layers, covering components from just above the hardware, up to the application itself. By offering an Elastic Server Community Edition to the global developer community, CFT tracks timely insight and trends about the evolving virtualized server environment at large, software component interaction information and proven configuration rules.

CohesiveFT's built-to-order application stacks:

- deploy to all major virtualization formats and an increasing number of cloud computing providers
- are given a unique identity, encapsulated, and injected with management and integration services
- include a connection to an online patch management system
- are able to track internal configuration changes for systems management and compliance purposes

The Elastic Server platform performs as a “content management system” for component assemblies allowing customers to ensure that only the components they want and approve go into their servers. It provides application stack lifecycle management and rapid provisioning; also recording each server's unique history, enabling monitoring and patch management. In the data center, unique software servers can be managed just like other network devices - each manageable, portable and maintainable. The speed and agility afforded by dynamic Elastic Server assembly allows organizations to reclaim resources spent on infrastructure maintenance, allocating them instead to new initiatives that will differentiate their businesses.

Still here? You probably want to one-click download the 7-page paper in its entirety. It's a compelling read, replete with fancy charts.



Cloud Cover: 2008 - 2009

Exclusive Viewpoint: History and Transformation of the Cloud

Edmund J. Sutcliffe

Thinker, OpenNovum



In their early days computers were tended by a small group of high priests who cared and tended them through the long dark tea times of their software cycles, and through the ups and downs of hardware failures.

As the years progressed a cult developed around building better solutions and over time these chapters got lost within larger churches of corporations. Each of these corporate churches presenting their individual ways to deliver their solutions, supported by their own particular belief patterns and approaches.

With the onset of cloud systems there is a new transformation of the working environment where the work of the many is replaced by the automation and consistency of the few. While the size and quantity of machines has changed, the manner of approach, the means of adoration, has returned full circle almost to where we were at the commencement of computing. Our “high priests” are returning, as experts in automation.

The increase in flexible working made possible by cloud computing has transformed the organisation's focus from attempting to control individual devices to monitoring how users are accessing the organisation's data. The organisation's entire infrastructure can now be controlled instantaneously from a single, centralised location. This means growing a business becomes much simpler as end-users can be added to systems quickly, no matter what their location. With such minimal setup costs businesses can afford to be bolder when expanding into new markets. This can be a beautiful tool for improving efficiencies, but there are disadvantages too. Disadvantages that are familiar to those of us who grew up with mainframe computing.

When writing for a mainframe, you spent lots of time thinking and little time running. As these early machines were accessed over networks and remotely by small control jobs and data sets delivered over slow remote connections, things were tested often by hand before they were released to run.

With the arrival of desktop computing sysadmins could get away with 'bodging' individual systems. As every system was different, serving individual purposes, it didn't matter that every sysadmin and every machine was singing from a different hymn book. Sometimes this

was even beneficial. We worked with trial and error and leaps of faith, rather than assessing the entire model before acting.

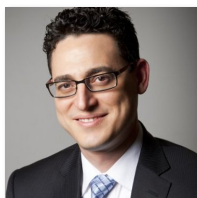
Now with the advent of Cloud, the assessment of the impact of a simple change upon the complex coupled systems becomes a difficult problem. If we get this assessment wrong, we are faced with problems as huge as the loss of air traffic control or an entire network of retail banking ATMs. The reputational effect of getting that wrong, the loss of faith, in fact, can be devastating.

But let us not forget that we are talking about engineering, not religion. We need to stop reinventing the wheel and use the body of knowledge we've already developed, through systems theory and operational research. The Apollo programme didn't reach the moon through trial and error, or through prayer. Yes, it *is* rocket science. At this point we no longer need the armies of repairmen and technicians we have been used to relying on in the past. A 'good' sysadmin is no longer good enough. The risks of applying any changes by hand are too great. People are poor robots. What we are looking for now is just a few visionaries, capable of understanding the entire model and applying the complex knowledge we already have. Not just clerics, but scientific saints.

Exclusive Viewpoint: Using CohesiveFT and AWS to Prove Cloud Concepts

Chad Lawler

Director, Consulting Services, Cloud Computing at Hitachi Consulting



When I first started working with CohesiveFT I was evaluating cloud networking tools for use with public cloud environments. In my role at the time in 2008, I was co-leading on a disaster recover (DR) proof of concept engagement for a major US-based retailer to explore the viability of public cloud for DR and failover scenarios.

As part of the 4-month POC, we determined the target viable application portfolio representative of these applications they would use in the cloud. Like many enterprise organizations, they had a vast array of application types and technologies. We narrowed the portfolio down to a Java web application, a Microsoft web application, and Websphere web application, which served as the representative use cases to test DR in a cloud environment. In diverse client environments, IT teams often think they have a good handle on their IT resources, server inventories, integration points, where applications live, and so on. However, once we got into the details of an application portfolio assessment, there were a few resources they didn't know about! As a result, we conducted a broad and detailed assessment to map out the applications along with all of the detailed components required to completely duplicate the applications in an alternative cloud environment. Working with numerous tiered and diverse IT teams, this portfolio assessment and mapping process took a month, just for three unrelated apps. Once complete, we began building the plan to duplicate these applications in the Amazon Web Services (AWS) public cloud.

At the time, AWS was far less mature than it is now. They did not have the EC2 admin UI other than command line. There was no AWS web console then, nor were there many of the services they have now like auto-scaling, load balancing, and platform VPN services. We knew we needed to connect the client's network with the Amazon network. We knew we needed to fail services over and back. We needed to connect networks and control the network route from the edge of AWS.

We looked around the market and found VPN-Cubed product (now called VNS3). We had talked to CohesiveFT previously, and this would be the perfect use case for collaboration between our teams. We used VPN-Cubed to create the virtual network between AWS and the

identified applications. For networking, there were no other real alternatives at the time, as Amazon was still evaluating enterprise needs.

We then created a duplicate of the three production apps: the Microsoft app, the Linux Java app, and the WebSphere web-services app. The only way to truly test DR is to failover fully from the production environment. You take production apps in production hours, and fail over. The risks associated with that are jaw dropping for most management team. Ultimately, our customer didn't want the risk for only a POC, so we created a demo environment to do the test and prove the failover.

In our POC, we didn't actually interfere with the client's production environment, so technically this POC was not a completely real-life disaster scenario test of production failover. However, we did duplicate the client's environments in their own-site DR lab, as well as in the AWS EC2 (Elastic Cloud Compute) environment, and executed complete failover and fail-back testing between the two.

In a nutshell, the POC was a tremendous success. With some process improvements, we ended up with a Recovery Time Objective (RTO) of 82 minutes. We were able to prove that DR is a fantastic use case for public cloud. Your DR footprint can be very small compared to production because you can have minimal environment systems in the cloud, say with just one or two server footprints per app tier. You can scale those up during DR testing, which is a critical must-have for any DR plan, but you only have to incur the full production footprint and scale costs during DR testing windows. DR in the cloud provided a very minimal cost-scenario for the client that was quite attractive.

Working with CohesiveFT was great. They were always available, always ready to pitch in and help, and had a great technical team who were a pleasure to work with. As Enterprise Architects leading this POC, we did not have a network engineer at our disposal and it had been a while since I had logged into a router command console. The CohesiveFT team handed the pre-configuration and routing setup we required. Their network engineers were very sharp and knew instantly how we should best configure our environments for the appropriate network failover and fail-back. The VPN-Cubed product and the CohesiveFT team handled most of the complexity for us, which was great.

Some lessons learned focused on determining actual customer needs, appropriate application fit in cloud, and how to handle the immaturity of the cloud offerings at the time. Most enterprise environments are very complex and enterprise applications are often unique in

configurations and implementations. It's of the utmost importance to fully understand and map out your complete application environments, including all of the integration points, in order to successfully recreate this in the target cloud environment.

We also learned that not every application is a good fit for public cloud, from a performance perspective or from a financial perspective. For example, DR is a fantastic use case for public cloud because you need only to run a minimal environment until DR testing or actual DR is required. On the other hand, if you're going to run an application 24/7/365 in an "always-on" scenario, you have to ask "what would it cost to run this in my data center versus the public cloud for the next 2 or 3 years?" In any of the major Cloud Service Provider environments, the costs can really start to add up.

The financial cloud TCO question can pose a problem, as most customers simply don't know how to determine the true, actual costs of their data center applications. There aren't always good chargeback models, and there are more costs in administration, manual effort, and other forms of human capital. It's very, very difficult for companies to quantify what they're spending on a per-application basis. As the saying goes "IT departments know their costs, but they don't know always know their economics."

I see companies struggling with how to adapt to cloud today. Initially, I thought it would be a mix of industries and clients looking to adopt cloud, but in fact many of my prospective clients are in fact product companies, traditional hosting providers and telecom companies trying to evolve to become enterprise cloud service providers. They are under incredible pressure to adapt, evolve and stay relevant. The products and services they now offer are not going to be competitive enough in the face of emerging cloud competition. I'm using the knowledge gained from this early enterprise POC to this day to provide insight to these service providers as they evolve and build their cloud services and solutions.

Thursday, October 16, 2008

We're #1, (and #2, kind of #3, and #6, and part of #10)

Patrick Kerpan - CohesiveFT Blog



Yesterday, Gartner Group announced their top ten technology trends for 2009.

Heading up the list of 2009 trends is Virtualization, followed closely by Cloud Computing. Since we started down the Virtualization and Cloud path more than two years ago, some might say that we saw this coming. We've

specifically architected our Elastic Server platform to meet these changes in the virtual landscape. Our automated "factory" lets you assemble custom virtual servers in minutes, and deploy to the cloud or virtual environment.

We are pretty excited about how big virtualization has become (thank you VMware for the 10+-year-hard slog), and how interesting the skyscape of clouds is becoming.

We started CohesiveFT with the premise that 4 major trends were colliding in the data center and were going to reshape the IT environment from development thru deployment; **open source, open standards, virtualization, and loosely couple architectures**. In many ways I think of cloud computing as the "first derivative" of these trends - and has become a trend of its own, a 5th driving trend reshaping the IT landscape.

Also, surprising was Gartner's #3 (Servers-Beyond Blades), #6 (Specialized Systems), and the now ever-present Green IT at #10. It seems like these are additional byproducts of the same forces. The way I see it, we have been fortunate to live through a "silent earthquake" in IT, and the aftershocks keep delivering smaller - but interesting enterprise impact.

If you're thinking about how you might adapt your IT department to meet these 2009 trends, come check us out. In the meantime, we'll keep doing the voodoo we do to stay ahead of the curve.

Competitive Landscape					
Cloud Security	Cloud "Brokerage" Migration and Integration	Topology / Appliance As A Service	Hypervisors	Management & Monitoring	DevOps
Cloudpassage Dome9 vCider Nicara	CohesiveFT Appcara RightScale rPath Verizon/ CloudSwitch ServiceMesh Racimi Kaavo	AppZero Jumpbox DotCloud Hateize Appistry GridRobotics CloudShare	VMware ESX Citrix Xen OSS Xen Parallels Hyper-V KVM Sun Open xVM	CA Unicenter Enstratus Ylastic HP OpenView IBM Tivoli Ylastic AccelOps MSFT MOM Zenoss Quest FluidOps Nagios Hyperic GroundWork TapInSystems	PuppetLabs Bitnami OpsCode IBM Icon DynamicOps ScaleXtreme Scalr Standing Cloud
HP Dell IBM Oracle 3leaf Lenovo Penguin Computing Cisco MorphLabs Marvell	Arrow Market Gale Tech. Layer7 CloudSoft Corp Fortisphere NewScale Gravitant Cordys	OpenStack Cloudscaling OpenStack Cloudia Scaleit CA-AppLogic Hexagrid onApp vCloud vSphere Citrix Delivery Center Plexiant Abiquo SkySys GoGrid Eucalyptus Globus Nimbus Cloudstack OpenNebula	Windows Server 2008 Google Chrome Debian Ubuntu CentOS Fedora RedHat Solaris Novell SLE openSUSE Connectiva Red Flag	Amazon AWS CloudSigma "vClouds" Zimory Terremark virtuStream Azure BT AppEngine Verizon Heroku OpSource Joyent GoGrid Aptana Rackspace	Apprio CSC Dell Perot Hitachi Consulting CapGemini Accenture IBM GTS IBM - like HP EDS Assimil8
HW Integrators / Assemblers	Cloud "Brokerage" Other	Public/Private Cloud Infra	OS Vendors	Public Clouds Managed Private	System Integrators

Tuesday, October 28, 2008

VPN Cubed – Cloud is Ready for the Enterprise

Krishnan Subramanian - Cloud Ave



CohesiveFT, a 3 year old Cloud Computing platform provider, today announced the release of an exciting new security product for the Clouds that could lure the enterprises into Cloud Computing. Their new product, called VPN-Cubed, provides a security perimeter for the IT infrastructure whether it is inside a single cloud or multiple cloud or hybrid cloud-datacenter

ecosystem. This is a pretty exciting product that has the potential to reshape the enterprise cloud adaption.

CohesiveFT was started in 2006 offering the Elastic Server Platform, an online service for configuring and assembling custom stacks to be deployed on the cloud infrastructure. Elastic



Servers are custom application stacks that can be assembled seamlessly using wide range of applications available. The granular control offered to the users allows them to just run the software they want and, in their words, nothing more nothing less. One can either use the pre-populated application stacks or stacks offered by other users or by building your own stack by cherry picking various components. Once the stack is assembled on the Elastic Server Platform, and can be deployed on any virtual environment.

The issue of vendor lock-in is one of the important issues facing companies planning to shift their operations to the cloud. According to Phil Clarke, Sales and Services Director at CohesiveFT, their platform offers an ideal solution to avoid such a vendor lock-in. He points out that any stack built using their elastic server platform can be deployed to any cloud infrastructure easily. It is true from the point of view of deployment but I think it doesn't help much in the migration of an existing stack from one cloud vendor to another.

Right now they have a free Edition suitable for developers and a very inexpensive Personal Edition for commercial purposes. They will be releasing a Team Edition sometime in the future satisfying the needs of bigger companies. Unlike Rightscale, which is involved in the management of Cloud instances after initial deployment, CohesiveFT focuses on helping their customers in assembling the application stack at a lower cost.

After developing a successful platform for the deployment of various application stacks in the cloud infrastructure, CohesiveFT is taking the next step in the evolution of a company by

tackling the security problem, the most important issue preventing the enterprise customers from joining the cloud computing bandwagon. Today's release of VPN-Cubed will offer a security perimeter covering the IT assets deployed in any kind of ecosystem (cloud or managed). One of the biggest worries for enterprise customers is the need to cede control of the security of their data to the cloud provider. With this solution, the company can retain the control of security even inside the cloud. VPN-Cube uses the popular Open Source VPN software, OpenVPN, to act as an encrypted LAN inside a single cloud or as an encrypted WAN across multiple clouds. This allows the cloud based clusters or the hybrid cloud-enterprise datacenter ecosystem to appear as a single physical network similar to Enterprise 1.0 infrastructure. VPN-Cube enables enterprise software requiring multicast networking and, also, offers greater control over network addressing, thereby offering an enterprise level secure network either inside a single cloud or across multiple clouds. With VPN-Cubed, companies can easily scale their infrastructure in the cloud with the ability to seamlessly add the newly provisioned resources into their existing VPN-Cubed network. Unlike the traditional VPN approaches commonly used to connect remote employees and offices, CohesiveFT's VPN-Cubed product runs inside the cloud-computing facilities with failover capability that keeps cloud clusters running reliably even in the event of a VPN server failure.

With the release of VPN-Cubed product, CohesiveFT now offers both an automated deployment option inside the cloud and a way to maintain control of the security. This makes cloud computing an attractive option for enterprises worried about the privacy and security of their data. This product will also help companies interested in having a hybrid cloud-datacenter ecosystem to protect their critical data by keeping on their servers while offloading the non-critical applications and data to the cloud. Right now, VPN-Cubed supports Amazon EC2, Flexiscale, Go Grid, Mosso and a few others. I spoke to Patrick Kerpan, CTO at CohesiveFT, and asked him if he expects enterprises to embrace cloud computing rapidly after the release of this much needed security tool, he was very optimistic about it and feels that the current economic downturn might force companies to consider cloud computing as a way to cut costs. Irrespective of whether we see a rapid enterprise adaption of cloud computing or not, I am pretty confident that VPN-Cubed will help companies control the security of their data.

The post was originally published by Krishnan Subramanian on CloudAve: <http://www.cloudave.com/2762/vpn-cubed-cloud-is-ready-for-the-enterprise/>

November 14, 2008

CohesiveFT VPN-Cubed: Not Your Daddy's Encrypted Tunnel

Chris Hoff - Rational Survivability



I had a briefing with Patrick Kerpan and Craig Heimark of CohesiveFT last week in response to some comments that Craig had left on my blog post regarding PCI compliance in the Cloud, here.

I was intrigued, albeit skeptically, with how CohesiveFT was positioning the use of VPNs within the cloud and differentiating their solution from the very well established IPSec and SSL VPN capabilities we all know and love. What's so special about tunnels across the Intertubes? Been there, done that, bought the T-Shirt, right?

So I asked the questions...

I have to say that unlike many other companies rushing to associate their brand by rubber cementing the word "cloud" and "security" onto their product names and marketing materials, CohesiveFT spent considerable time upfront describing what they do not do so as to effectively establish what they are and what they do. I'll let one of their slides do the talking:

CohesiveFT Update

- We are not a cloud. We are a virtualization and a cloud-computing complement
- Our automated Elastic Server® factory helps individual developers and small teams assemble individual servers and deploy to the cloud or virtual environment.
- Enterprise customers require clusters of servers, configured to fit their unique infrastructure and use-case
- Security is the gating factor preventing Enterprise cloud adoption
- The next logical step on our path to enablement is to remove this barrier and help Enterprise users confidently leverage the cloud
- Our new packaged service, VPN-Cubed™, is a complement to our automated platform

cohesiveft.com/vpncubed 

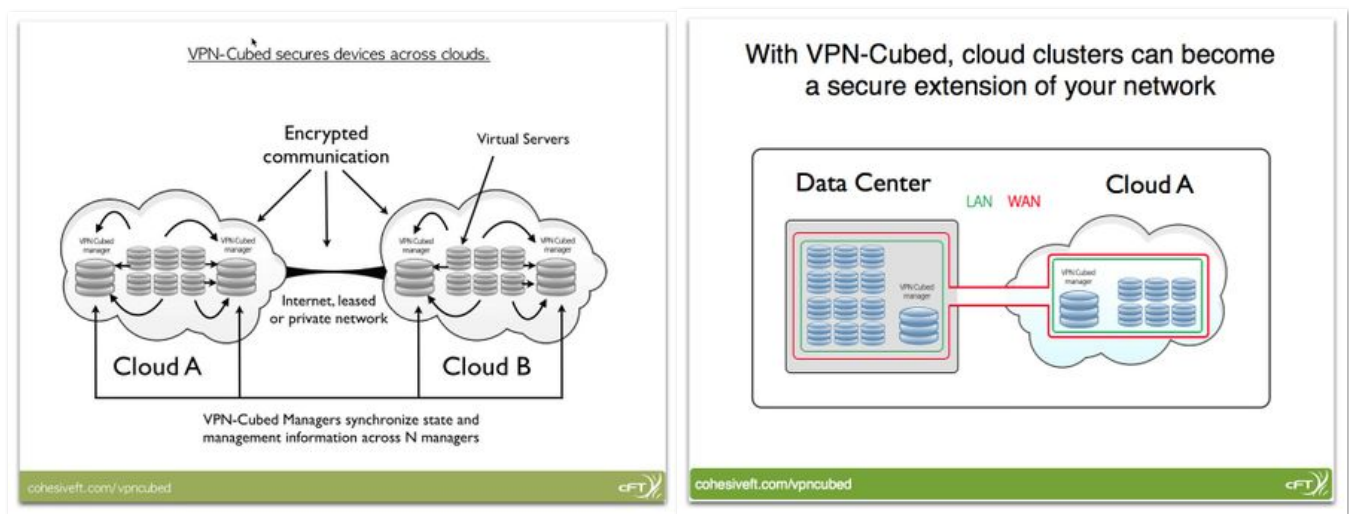
Okay, so they're not a cloud, but they provide cloud and virtualization services...and VPN-Cubed provides some layer of security along with their products and services...check. But...

Digging a little deeper, I still sought to understand why I couldn't just stand up an IPSec tunnel from my corporate datacenter to one or more cloud providers where my assets and data were hosted. I asked for two things: (1) A couple of sentences summarizing their elevator pitch for VPN-Cubed and (2) a visual representation of what this might look like.

Here's what I got as an answer to question (1):

"VPN-Cubed is a novel implementation of VPN concepts which provides a customer controlled security perimeter within a third-party controlled (cloud, utility infra, hosting provider) computing facility, across multiple third-party controlled computing facilities, or for use in private infrastructure. It enables customer control of their infrastructure topology by allowing control of the network addressing scheme, encrypted communications, and the use of what might normally be unrouteable protocols such as UDP Multicast."

Here are two great visuals to address question (2):



So the differences between a typical VPN and VPN-Cubed comes down to being able to securely extend your "internal ~~clouds~~ infrastructure" in your datacenters (gasp! I said it) in a highly-available manner to include your externally hosted assets which in turn run on infrastructure you don't own. You can't stand up an ASA or Neoteris box on a network you can't get to. The VPN-Cubed Managers are VM's/VA's that run as a complement to your virtual servers/machines hosted by your choice of one or multiple cloud providers.

They become the highly-available, multiprotocol arbiters of access via standardized IPSec protocols but do so in a way that addresses the dynamic capabilities of the cloud which includes service governance, load, and "cloudbursting" failover between clouds -- in a transparent manner.

Essentially you get secure access to your critical assets utilizing an infrastructure independent solution, extending the VLAN segmentation, isolation and security capabilities your providers may put in place while also controlling your address spaces within the cloudspaces

encompassing your VM's "behind" the VPN Managers.

VPN-Cubed is really a prime example of the collision space of dynamic application/service delivery, virtualization, security and cloud capacity governance. It speaks a lot to re-perimeterization that I've been yapping about for quite some time and hints at Greg Ness' Infrastructure 2.0 meme.

Currently VPN-Cubed is bundled as a package which includes both product and services and supports the majority of market-leading virtualization formats, operating systems and cloud providers such as Amazon EC2, Flexiscale, GoGrid, Mosso, etc.

It's a very slick offering.

/Hoff

The post was originally published by Chris Hoff on Rational Survivability: <http://www.rationalsurvivability.com/blog/2008/11/cohesiveft-vpn-cubed-not-your-daddys-encrypted-tunnel/>

December 10, 2008

2009 in Virtualization and Cloud Computing: The Year of the Virtualization Professional

Patrick Kerpan - VMBlog



Virtualization is the key architectural break that is driving a massive transformation in the world of computing, a transformation that will affect the way businesses work and

people live. It is computing “gone through the looking glass”. Ten years from now we will look back and say “how did we used to do that before...”, just like we do when reflecting on personal computing or the Internet.

When, exactly, did computing, like Alice in Wonderland, press against the mirror and fall through? My guess is sometime after VMware began driving this modern wave of virtualization (I know, I know, IBM mainframes had “virtualization”) and sometime before this year.

Think back to the “relatively” simple trick of virtual memory. Many of you remember, those of you younger can look it up on Wikipedia. We went from counting bytes and manipulating overlay files to virtually unlimited memory. This transformed how we built software systems. It changed the skill set of whole sections of the programming profession. People who couldn't get used to “wasting” memory potentially slowed projects down, despite the fact that a year or two before, their special skills in memory paucity made them saviors.

Just as virtual memory meant nearly unlimited memory for your application – virtualization means a future of nearly an unlimited number of computing devices. But that's a few years off. What does this mean for 2009? Well – 2008 was the year that began to prove the model; leading vendors like VMware, Citrix and Microsoft showed continued and clear commitment to the market, the first wave of virtualization startups that didn't survive provided some clarity in the shape of the market structure, trade press and analysts really started to pay attention, and the powerful complementary trend of cloud computing emerged as a “top of mind” topic. In short, 2009 is the year things begin to “get real”.

This coming year brings two markets to deal with; customers with “cloud intent”, working diligently to derive value from virtualized infrastructure now, and the beginnings of the traditional enterprise market, always a 10 year acculturation to concepts, vendors, and products. Those customers with “cloud intent” (comprised of a mix of Web 2.0 companies,

SMB, and Enterprises) are the most likely to move whole parts of discrete infrastructure to cloud computing centers (Amazon EC2, Flexiscale, GoGrid, ElasticHosts, etc.). Traditional enterprise will want to have “hybrid clouds”; virtual infrastructure in their private data centers connected to cloud facilities for “fill in the blank” reasons (experimentation, high watermark, overflow, seasonal, business unit politics, and more).

Another 2009, “getting real” data point is Gartner Group's top 10 technologies for '09. In case you didn't see it Virtualization is #1 and Cloud Computing is #2. But if you just stopped there, look further for the impact of virtualization. “Servers Beyond Blades” (#3) is about hardware becoming more dynamic. “Specialized Systems” (#6) hearkens to the emergence of deployable clusters of technologies – well beyond the simple virtual appliance in sophistication and power. And of course, 2008's #1 is this coming year's #10, Green IT.

So here we are on the cusp of '09. The virtualization vendors continue to extend and differentiate as much as they standardize, new clouds announce almost weekly, SaaS is a foregone conclusion as a strategy for the enterprise, PaaS is growing in popularity, the OS vendors are creating virtualization tuned distributions, and open source as the source of corporate middleware increases in scope. Each of these market solutions has the potential to create economic opportunity for the enterprise. When combined, the skills to leverage them isn't the same old job. Enterprise IT staff at the forefront as users and buyers of this confluence are a new type of IT professional, the virtualization and cloud computing professional.

Do the math. What is the ubiquity of virtualization? Pick your number. How about Virtual Machines per server? You pick. Refresh rate per year. You pick. It doesn't take a lot of arithmetic to see that we are going to end up with multiples of an order of magnitude of new computing device every year compared to the not too distant past. As an industry if we are assembling, using, and discarding 10 million VMs per year, 20 million, 50 million, 250 million maybe by 2015, then some things have to change.

Emergency? Opportunity? Probably a bit of both. But, then again, the year is only 2009. We are just starting to get real. You have a couple of years to get ready for this new reality, well actually, virtuality.

Originally published in the VMblog Virtualization and Cloud Prediction Series. For the original post, visit(<http://vmblog.com/archive/2008/12/10/2009-in-virtualization-and-cloud-computing-the-year-of-the-virtualization-professional.aspx>) All rights reserved by the author.



Fog Surrounds Cloud: 2009 - 2012

Exclusive Viewpoint: Cloud Computing: Past, Present and Future

Krishnan Subramanian

Founder, Rishidot Research



After Amazon Web Services (AWS) launched cloud services in 2006, the adoption has been steadily increasing with more and more enterprises embracing cloud today after overcoming their initial fears on cloud security. The initial idea of cloud, mainly defined by the public cloud services, pushed the industry to rethink IT operations and application development. Early success of SaaS combined with the popularity of mobile devices in the enterprise, paved the way for large scale adoption of cloud computing. Today, enterprise IT managers are not wrestling with the question of whether they should adopt cloud but, rather, asking how they can maximize the benefits offered by cloud computing.

The initial developer enthusiasm for cloud computing can be traced to the new application architectures thrust upon developers by public clouds and the ease with which they could take advantage of agile development methodologies. Even though developer experience was a big attraction towards public cloud services, many developers also saw the opportunity unleashed by cloud computing by lowering the barriers to entry. Cloud computing helped the smart minds in the industry take their idea to market much faster than any time in the past at an affordable cost. Now a startup, or even an individual developer coding in the basement of his/her house, could have access to computing resources previously available for large organizations with tons of money. This ushered in a new era of entrepreneurship leading to services that were previously dreamed by science fiction writers. The rate of innovation has accelerated so much in the industry that newer tools and services were made available to consumers at a daily basis.

Cloud computing not just transformed the lives of the citizens of advanced nation but it also opened up opportunities for countries who were previously not part of the global economy. It offered an easy ramp for entrepreneurs in poor or war torn nations to take their goods and services to the global market. Cloud not only helped entrepreneurs from different parts of the world to compete in the global economy, it also opened up world class education to be available to billions of people around the world. For the first time in the history of human beings, there was some hope to lower the economic divide between the wealthier nations and

the poor. Soon experts saw potential for cloud computing in many verticals affecting the global economy such as the financial sector, health care, manufacturing, etc.

As more and more startups started to take advantage of cloud services, individual developers and smaller teams inside enterprises bypassed their IT to take advantage of cloud services offered by companies such as Amazon, Google, Heroku, etc. They realize that cloud computing offers them an easy path to gain compute resources at a cost that can easily be expensed. This trend helped individuals and teams inside the enterprise innovate rapidly without worrying too much on the high cost of failure. Even though cloud computing empowered developers and business users inside enterprises, the rogue IT caused considerable headache to enterprise IT. They were deeply concerned about potential security and governance problems these cloud services can create for their organizations. Even though many of the cloud services in the market were suitable for enterprise adoption, IT departments were not ready to embrace them. The lack of visibility and control were a big inhibitor for most of the enterprises. Moreover, enterprises had invested heavily on legacy workloads that were not suitable for use in public clouds. Enterprise IT understood the importance of cloud services for their business but they were not ready to trust public cloud services.

Such a predicament is a great opportunity for industry willing to experiment and innovate. Vendors saw an opportunity to offer cloud platforms that could meet the needs of enterprise IT while also empowering their developers and business users. Soon the market realized that cloud platforms that could help enterprise IT offer cloud services to their end users while also maintaining a clear control over the compute resources is a huge opportunity which could potentially change the enterprise IT landscape completely. Companies like VMware and open source projects like OpenStack, CloudStack, and Eucalyptus jumped in to meet this enterprise need for private and hybrid clouds. Even though private cloud appeared to take off big in the enterprise market, the benefits of public clouds were too good to be ignored by the enterprise. This lead to another iteration of cloud computing and today hybrid cloud is a reality.

The benefits of public cloud like self service, elasticity, metered billing coupled with enterprise grade features such as hardened security, governance, etc. are the key drivers of hybrid cloud in the industry today. There is considerable debate on where the industry is headed among the vendors, pundits, etc. Traditional vendors who are offering private cloud platforms believe that most of the workloads will remain in the private cloud while some of the low risk workloads move to the public clouds. Progressive vendors like Amazon, Google and others

believe that even though most of the workloads are on private clouds today, future innovations will bring most of the workloads to the public cloud services eventually leading to the death of private clouds. Even though there is no clear consensus on the future of enterprise cloud journey, there is no going back from the services based world and cloud computing is here to stay.

As enterprises get comfortable with cloud and plot their hybrid cloud strategy, there is a quiet trend emerging from the sidelines where platform services are slowly gaining traction. The containers-based approach taken by these platform services are very attractive to enterprises because the interfaces offered by these platforms empower the developers to be more productive while making the IT operations less complex. In the future, we will move from thinking in terms of Virtual Machine chunks to Application Containers. Enterprise IT will be focussed on the orchestration, automation, management and governance of these applications containers. Application containers makes service delivery and consumption much more seamless and platform services not only will make developers productive and innovate more rapidly, it will also make it easy for business users to take advantage of these services for higher order innovations. In the not so distant future, innovation in the enterprise market will be on par with the consumer sector and it will be driven by cloud services underneath. It is time to tighten your seat belts and get ready for a journey which will be completely different from the past.

Exclusive Viewpoint: White Knight Solutions from CohesiveFT

James Elwood

CTO, Geezeo



In financial services, banks sometimes farm out parts of their architecture to various firms. For Geezeo, we started business knowing we wanted to build cloud-only software as a service applications for these financial institutions. The security needs of the financial services industry and our partners meant there was just no way we could use public internet to secure transaction data.

We knew our SaaS products and platforms were solid, but we were faced with an huge networking issue as each new partner signed: how do we secure this data without using public internet or building out middleware?

In early 2009, we started trying to build it on our own. We were trying to figure out how to build our own solution, knowing that our one new engagement might turn into 10 within a year. We'd cobbled together a little bit of IPsec, a little bit of OpenVPN and a whole lot of janky SSH tunnels. But what we'd designed originally just wasn't going to scale. The end solution we needed was a tunnel, so we looked into VPN services. After probably 2 months of tearing our hair out, we found what was then VPN-Cubed. It really was a "white knight." We dropped the product in, and within 2 weeks we had the first tunnels up and running.

VNS3 is all we've used to date. The wildcard is that we're still purely in Amazon AWS EC2. Without the physical infrastructure, being able to set up multiple IPsec tunnels to various data centers and the ability to keep them separate would have been an ongoing management challenge. You still can't do that with Amazon, even with their VPN offering. The beauty of the solution with CohesiveFT is that VNS3 was able to become pretty much a giant edge kit.

The biggest selling point for us is the UI. A single sysadmin manages 30 tunnels easily at one time, as if it was a Cisco appliance. It would have been a mess in Amazon or it would have to take a lot of custom work. Now we bridge data centers into specific assets of our infrastructure without having to expose the whole back end. Another beauty of the solution is that it's been rock solid and it's scaled very nicely. We went from maybe 4 customers in 2009 to over 240 now, and plans to grow internationally. To be honest, we will probably use VNS3 as a mainstay of our architecture when we grow. We might even end up leveraging it internally for things.

Friday, August 28, 2009

New VPN-Cubed Version and The Cloud Connectivity Market

Ryan Koop - CohesiveFT Blog



The people we encounter looking to leverage cloud computing come in all shapes and sizes. Some are large enterprises with massive deployments, others are just looking to deploy a front end and one database. Clearly one size doesn't fit all, so why should we only offer one IPsec-to-cloud solution? Our

Free Edition is a direct response to our users looking for a less costly solution to join the movement to cloud security and control. Those users now have an excellent opportunity to start extending operations into the public cloud without a large initial cash outlay. (Rule 76: No excuses, cloud compute like a champion.)

The Cloud Connectivity Market is Now Open

Cloud vendors, in this instance Amazon Web Services, have realized there is not going to be wholesale migration of data centers to clouds. The recent release of Amazon Web Service's Virtual Private Cloud (VPC)

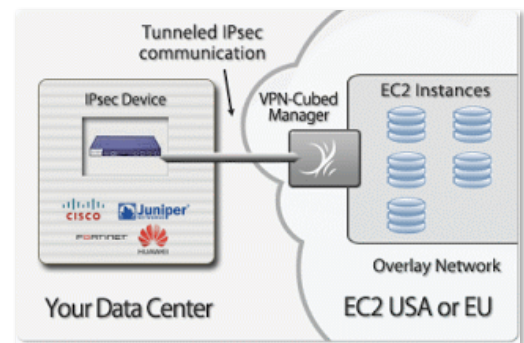
helps to move the conversation of security and control in the cloud forward. Having a major player join us in the ever evolving Cloud connectivity market is extremely exciting! There is no better validation of our original idea, VPN-Cubed, than a complementary offering from AWS. So in the spirit of friendly competition, we would like to explore some of the similarities and differences between our two offerings.

While the AWS VPC is somewhat similar in basic approach and feature set to our VPN-Cubed offering, Amazon's VPC beta does not target key areas of customer control and market interoperability (see our comparison matrix below).

What does it all mean?

Clearly, we will overlap in some segment of the market, but there is a major difference in our approaches. We are a software company and have created VPN-Cubed using software virtual appliances. AWS are raised-floor gurus, they use hardware expertise to provide connectivity solutions.

With out knowing much more about Amazon's implementation we won't make any assumptions about future capabilities and limitations. But it is safe to assume AWS VPC will



emerge from beta with a broader feature set than it currently enjoys. Below is a "what's available today" comparison. Enjoy it.

Again, kudos to the Amazon team for realizing and addressing the key adoption hurdles for cloud computing: security and control. We look forward to integrating our VPN-Cubed Manager into the VPC infrastructure and collaborating to bring the most comprehensive set of features to cloud bound customers. Two is a crowd.

AWS VPC vs. CohesiveFT VNS3

Feature		AWS	VNS3 Enhances	VNS3 Extends
AWS Interoperability	Features available in all zones of EC2 USA today	✓	✓✓	
	Features available in all zones of EC2 EU today	✓	✓✓	
	Features integrated to EC2 existing security lattice (EC2 Security groups)	✓	✓✓	
	Can use EC2 Elastic IP Addresses	✓	✓✓	
	Ability to use Amazon load balancing service today	✓	✓✓	
	Access to Amazon S3	✓	✓✓	
	Support all EC2 Instance Types in All Regions and Zones	✓	✓✓	
	Ability to use Elastic Load Balancers across VPCs within a region			✓✓✓
	Ability to use Elastic Load Balancers across VPCs across regions for failover			✓✓✓
	Ability to use Elastic Load Balancers across VPCs within a region			✓✓✓
	Ability to use Elastic Load Balancers across VPCs across regions for failover			✓✓✓

Feature		AWS	VNS3 Enhances	VNS3 Extends
AWS Availability	Multiple VPCs per AWS Account	*	✓✓	
	Multiple VPN Gateways per AWS Account	*	✓✓	
	Multiple Customer Gateways per AWS Account	*	✓✓	
	Multiple VPN Connections per VPN Gateway		✓✓	
Address Control	Can ASSIGN SPECIFIC addresses to specific servers in my "VPC"	✓	✓✓	
	Create a Virtual Private Cloud on AWS's scalable infrastructure, and specify its private IP address range from any block you choose.	✓	✓✓	
	Divide your VPC's private IP address range into one or more subnets in a manner convenient for managing applications and services you run in your VPC.	✓	✓✓	
	Private IP Address Range Shared across Multiple Clouds and/or Virtual Infrastructures			✓✓✓

Feature		AWS	VNS3 Enhances	VNS3 Extends
Protocol Control	Allow customers to use BGP	✓	✓✓	
	Can use UDP multicast in my EC2 subnets			✓✓✓
	Can use UDP multicast between EC2 regions			✓✓✓
	SSL VPN Support			✓✓✓
	Multicast between data center and EC2			✓✓✓
	Support GRE Termination			✓✓✓
	Custom Layer 3 protocol modules (services based)			✓✓✓
Topology Control	Traffic can be routed directly to the Internet and NOT back across the Internet, into my datacenter and back out again	✓	✓✓	
	Securely route traffic to EC2 EU from EC2 US without having to route through the datacenter	✓	✓✓	
	Custom topologies & design services (declarative topology description)	✓	✓✓	
	Provides outbound NATing from Private VPC subnets	✓	✓✓	
	End user VPN Clients can connect to VPC using SSL Client			✓✓✓
	End user VPN Clients can connect to VPC using IPsec Client			✓✓✓
	Dynamic route updates available to SSL and IPsec Clients			✓✓✓
	Ability to move IP addresses between virtual infrastructures or clouds			✓✓✓

AWS VPC vs. CohesiveFT VNS3

Topology Control (cont'd)

Security Control

Feature	AWS	VNS3 Enhances	VNS3 Extends
Ability to create Cloud-based WANs that integrate corporate sites, cloud infras, partner sites, and colo or MSP infra.			✓✓✓
Provides outbound NATing from PublicVPC subnets			✓✓✓
Allows port forwarding from Internet to select inside VPC servers			✓✓✓
Route traffic between your VPC and the Internet over the VPN connection so that it can be examined by your existing security and networking assets before heading to the public Internet.	✓	✓✓	
Control inbound and outbound access to and from individual subnets using network access control lists.	✓	✓✓	
Bridge together your VPC and your IT infrastructure via an encrypted IPSEC connection.	✓	✓✓	
Network firewall controlling the VLAN	✓	✓✓	
Intrusion/Extrusion detection in the cloud - monitoring x-cloud subnets			✓✓✓
Access controlled on the host level by a unique cryptographic credential per virtual network address.			✓✓✓
Cryptographic identity linking (and segregating) multiple gateway routers			✓✓✓
Remote Support controlled by multi-organization (customer and vendor) 2-factor authentication			✓✓✓

Market Interoperability

Feature	AWS	VNS3 Enhances	VNS3 Extends
Windows and Linux device support	✓	✓✓	
Supports industry standard security appliances NAT'ed behind customer edge (Cisco ASA for example)	✓	✓✓	
Eucalyptus to EC2 support			✓✓✓
vCloud to EC2 support			✓✓✓
GoGrid/Rackspace/ElasticHosts/CloudSigma/Flexiant/etc - to EC2			✓✓✓
OpenStack to EC2			✓✓✓
IBM Smart Cloud and Smart Cloud Plus to EC2			✓✓✓
Easily integrate mobile phones and tablets to VPC infrastructure			✓✓✓
Citrix Virtual Infra to EC2			✓✓✓
Parallels Virtual Infra to EC2			✓✓✓
KVM Virtual Infra to EC2			✓✓✓
VMware Virtual Infra to EC2			✓✓✓
Let other AWS accounts (Partners, ISVs) launch instances to talk to VPC owner's instances directly			✓✓✓

Enterprise View

Feature	AWS	VNS3 Enhances	VNS3 Extends
2-way failover in VPC	✓	✓✓	
Instance can be both be part of a VPC and accessible to the general Internet	✓	✓✓	
Ability to create N-number of IDENTICAL defined subnets without routable connectivity allows significant gains in dev/test/staging.	✓	✓✓	
Web-based management interface	✓	✓✓	
Support for customer's IPsec endpoints behind NAT	?		✓✓✓
N-way failover in VPC			✓✓✓
Support for 3DES and AES 256 encryption			✓✓✓
Common abstraction model/interface across all clouds and virtual infrastructures			✓✓✓
Geographic or datacenter redundancy from customer side to VPC			✓✓✓
Emergency access possible if IPsec connection is down.			✓✓✓
Ability to connect a single VPC to multiple datacenters directly, as opposed to daisy-chaining datacenters via customer WAN.			✓✓✓
Ability to directly "dump" the interfaces to see traffic traversal and connection attempts.			✓✓✓
SNMP support for popular Enterprise monitoring systems.			✓✓✓

Dec 30 2009

2010 in Cloud Computing: GAME ON!

Patrick Kerpan - VMBlog



Looking back on my 2009 predictions they were "thematic". I think I did a decent job of capturing the spirit of the year - although I didn't particularly foresee

the "cloud washing" that was going to take place. I am not sure, but I think even my local dry cleaner was offering some sort of cloud solution in 2009. (As a note, at CohesiveFT we are liberal in our cloud definition, so when I say "cloud" I mean public, private, hybrid, etc..)

This year I will aim to be a bit more predictive - and with mere hours left on '09, here goes.

Enterprise adoption of cloud computing enabled solutions will begin to take off.

In 2010 we finally will begin to get beyond the early adopters. This is not the "early majority", but perhaps the beginning of the early majority. I call this the "early pragmatists".

Enterprise "Early Pragmatists" will begin initial deployments of agile infrastructure and "private cloud" based on Windows Server 2008 R2.

The reality of the enterprise is that Windows is the deep end of the pool. The whole agile infra and private cloud game has been playing out in the Linux world to date - not the enterprise sweet spot. Check your email inbox and the trade magazines, the Windows Server 2008 push is on by Microsoft, by its system integrator partners and the big resellers like Dell and HP.

This is the starting gun for the next generation enterprise data center.

"POA" (plain old applications) becomes the dominant cloud story.

Up until now the cloud story has been dominated by examples like Animoto and UK Channel4 - variable, huge scale stories. In some ways these successes have confused the enterprise - plain old IT guys hear these examples and think "I don't need that, my applications don't scale, don't need to scale". Enterprise IT lives out its life one business topology at a time. These topologies have 5 to 20 servers in them, fairly tightly coupled, and don't talk to other business topologies. This is "plain old applications". We did a project with one customer that has 140 such topologies, ranging between 4 and 7 servers, and broadly speaking none of them talk to each other. Helping enterprises embark on the long, slow migration of these systems to agile infrastructure is the payoff. Statistically this is where all the money is - and this is where the market will begin to focus.

Another public cloud super scale provider will NOT emerge.

For Infrastructure as a Service, Amazon is the only super scale player. The big market participants with billions in capital will announce more and more clouds and cloud services - but there won't be a super scale infrastructure that gives you easy, credit card based, no sales people, API-driven use of virtualized containers. My icon for this is HP. They will announce wondrous, magical things in press releases, complete with quotes from customers - but somehow the rest of us won't be able to see it or use it.

"Cures for baldness" and other magical claims begin to be called into question.

Ok, I understand positive spin, but outright lies are starting to bug me. As a "follically challenged" person, I would know if there was a cure for baldness, because there would be no bald people. Part of the cloud space intersects with systems management, automation and provisioning. I find this part of the market littered with uncertain cures to enterprise ills. One of the code phrases to look for should be "Now you can seamlessly provision into the cloud". With all of CohesiveFT's capabilities like multi-cloud image management, multi-cloud overlay network, and multi-cloud configuration management we certainly agree that is the GOAL of enterprise cloud computing. That said, I can guarantee you if any of the vendors that offer "seamless provisioning", "complete agentless discovery", or "automatic provisioning to any environment" actually could provide what they claim, then there would be NO SCREWED UP ENTERPRISE IT ENVIRONMENTS! I think analysts, trade press and customers need to have a more skeptical eye to claims like this and believe we will see some push back on magical claims in 2010.

Standards emergence will still be slow.

I still think we are in an early, sort of "pre cambrian" era of cloud computing and agile infrastructure. This isn't a statement against standards - it is just my view of where we still are, which is expanding as much as contracting. I do have a preference for "doers" proposing standards, instead of "talkers".

The hardware support for software encryption conversation begins.

Hybrid, public, and even some private cloud use-cases mean everything needs to be encrypted from end to end - always - by default - end of story. We need hardware assisted encryption on board. Intel, AMD, Dell, IBM talk among yourselves.

Originally published in the VMblog Virtualization and Cloud Prediction Series. For the original post, visit(A <http://vmblog.com/archive/2009/12/30/2010-in-cloud-computing-game-on.aspx>ll rights reserved by the author.

12 January 2010

Getting out of the weeds

Chris Swan - The State of Me



After some reflection on my recent series of posts about Paremus ServiceFabric on EC2 I realise that I never provided a high level commentary on what each of the moving parts does, and why they're important.

CohesiveFT Elastic Server on Demand – this is a factory for virtual appliances.

I used it to build the Amazon Machine Images (AMIs) that I needed. A bit like OSGi it uses a concept of bundles, and for some of the software that wasn't already there in the factory (e.g. the Paremus stuff) I had to create my own. Once I had the bundles that I needed I was then able to choose an OS, and build a server to my recipe (aka a 'bill of materials'). The factory would send me an email once a server was ready (and optionally deploy and start it for me straight away).

CohesiveFT VPNcubed – this was the overlay network that ensured that I had consistent network services (that supported multicast) covering the private and public pieces of the project. It basically consists of two parts:

A manager – which can exist in the private network or the cloud (or both). For simplicity I went with a pre packed AMI hosted on EC2

A set of clients. These are basic OpenVPN clients. For my AMIs I used a pre packed bundle. For the machines on my home network I just downloaded the latest version of OpenVPN. The manager provides 'client packs' containing certificates and configuration files, which need a little customisation to specify the manager location.

CohesiveFT ContextCubed – this provides the ability to start and customise a bunch of virtual appliances (AMIs) automatically. With the help of their CTO, Pat Kerpan, I was working with a pre release of this service (hence no link). ContextCubed (which I accidentally called ConfigCubed in my post about it) provides an init.d style mechanism that sits outside of the virtual machine itself. I used it to download and install VPNcubed client packs, start the VPN, stop some services I didn't want, reconfigure the firewall to allow multicast, and add binding config to the Paremus Atlas service (before starting it up). I could have also used it to create hosts files to work around some of the naming issues I encountered, but I think I'll wait for Pat to fix things up with DNScubed or whatever he ends up calling it. Hopefully in due course the *cubed services will all find their way onto the same virtual appliance, so there can be a one

stop shop for stuff that makes an application work in a hybrid cloud (or whatever suits your taste from private to public).

One thing that would have been fun to try (but that I didn't attempt) is closing the loop between the PaaS management layer in ServiceFabric, and the IaaS management layer in ContextCubed. This would allow (for instance) extra machines to be deployed dynamically to satisfy peaky workloads (or deal with failure) running on ServiceFabric. I'll leave that for another day.

The original post is available on Chris Swan's personal weblog: <http://blog.thestateofme.com/2010/01/12/getting-out-of-the-weeds/>

Tuesday, November 2, 2010

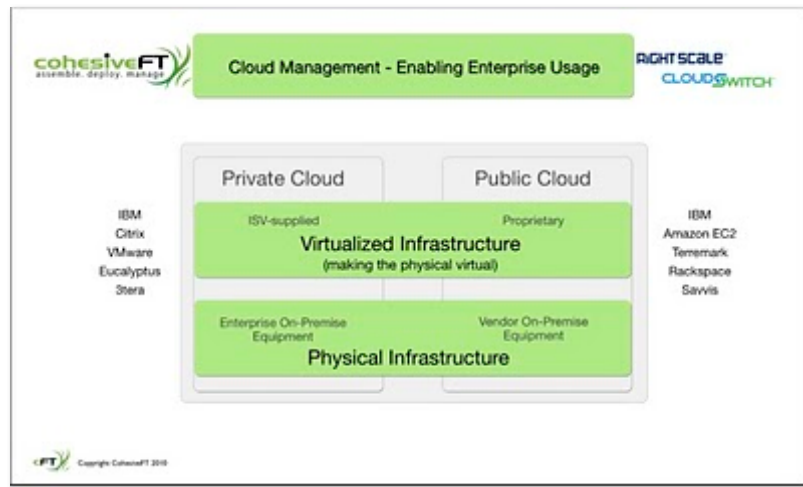
Welcome to the User-Cloud (Part 1)

Patrick Kerpan - CohesiveFT Blog



Disclaimer: I am talking about IaaS when I use the term "cloud" in this post.

Welcome to the world of cloud tenancy – where we (we being people who use the cloud as opposed to people who provide the cloud) are shared tenants of increasingly powerful virtual, agile, api-driven and anonymous compute and storage mechanisms. This is an area where I continue to see confusion in how people look at the world of IaaS (hereafter called “cloud” in this post). In the picture below is CohesiveFT's (CFT) simple model of the market structure where we have the physical layer (been around for ever), the virtual infrastructure layer (been around in growing forms for 10+ years, dominated by VMware), and the user-cloud or cloud-tenant layer (relatively new). By user-cloud, I don't mean the end user of the application, rather the person or team responsible for the application topology being run on the cloud.



The new "separation of concerns"

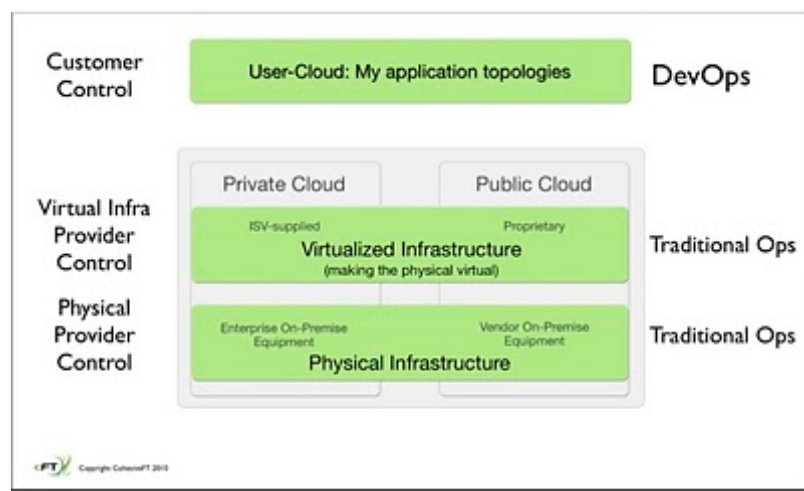
If you think about it, when you are using AWS – the team of people running the physical infrastructure have no direct interaction with your application topology; and you don't want them to.

The teams of people running your virtual infrastructure at AWS have ALMOST no interaction with your application topology; and you don't want them to. The virtual infra provider team might be making snapshots of your VM's and Storage as part of their cloud service. They might do some sort of LAN-based server motion where they move your VM's 10 to 50 feet

from Server A to Server B, if it appears Server A is about to fail. Other than that – you don't want their involvement – what could they do? They see your topology from “below”, they see a clump of x86 VM workloads, and have no knowledge of the semantics of your application topology, and in fact shouldn't.

You, in the user-cloud, know the semantics of your application, you know how it scales, you know where it came from, where it is going, and whether it can migrate from region to region or country to country. And to date you have used either your own self-developed knowledge of EC2, private Eucalyptus, Terremark, etc., or you used Rightscale or CohesiveFT, or one of the other tooling sets that fit in the category Gartner calls “Cloud Management Platform”.

I believe this is an industry wide trend, and basically creates a new “separation of concerns” in how IT functionality is provided both by vendors and IT staffs, and how IT functionality is consumed. The Enterprise IT environment has vacillated for decades between centralize, decentralize, centralize, etc.. This new separation of concerns allows for the gradual evolution of a hybrid IT organization where there is a logical, and purposeful bit of both models.



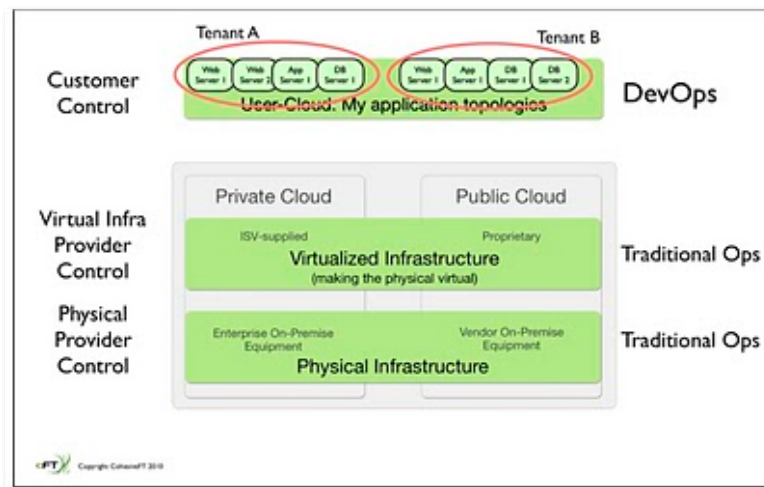
There are now three distinct control planes; physical, virtual, and customer / user.

Physical infra is a centralized activity, at the very least at the geographic level.

Virtual infra will be centralized at the virtual infra layer, with the newer versions of virtual infrastructure software offering some interesting geographical federation capabilities. It is possible one might be running both VMware and Xen virtual infrastructure globally, and those virtual infra's will likely be in the control of either one IT group, or one IT group per vendor solution.

Application topologies ideally run “when” (in time) and “where” (in geography) they are most useful to a business. This means detailed knowledge of the wants, needs, beliefs, and fears of the consuming business units. If this is a global customer database without a predominant business unit owner – then that application topology itself might be managed by a centralized “common applications” team.

If it is a cornerstone application for a company's Swiss Private Client group – then I am guessing it will be managed by a business line oriented IT group in Switzerland, likely to be running on a piece of Swiss-domiciled virtual infrastructure.



The key here is to realize that there are distinct products/vendors for each layer of your cloud infrastructure, and there are distinct buyers/users for each layer of your cloud infrastructure.

To me this sets up nicely to fit into the UK idiom of “horse for courses”.

*However, in part two of this topic I will address the problematic nature of the new separation of concerns. The gist of the issue is there is a temptation in the “enterprise cloud” or “private cloud” to violate them, not recognizing that there are now consumers of each level's vendor offerings, as well as a desire of course by vendors to be the single source of all cloud computing capabilities for an enterprise (not necessarily a bad thing) by providing one monolithic product that serves all levels (a bad thing).

(pat k)

Friday, November 5, 2010

Welcome to the User-Cloud (Part 2)

Patrick Kerpan - CohesiveFT Blog



Disclaimer: I am talking about IaaS when I use the term "cloud" in this post.

A quick recap from Part 1 where I was outlining what I see as:

The new "separation of concerns"

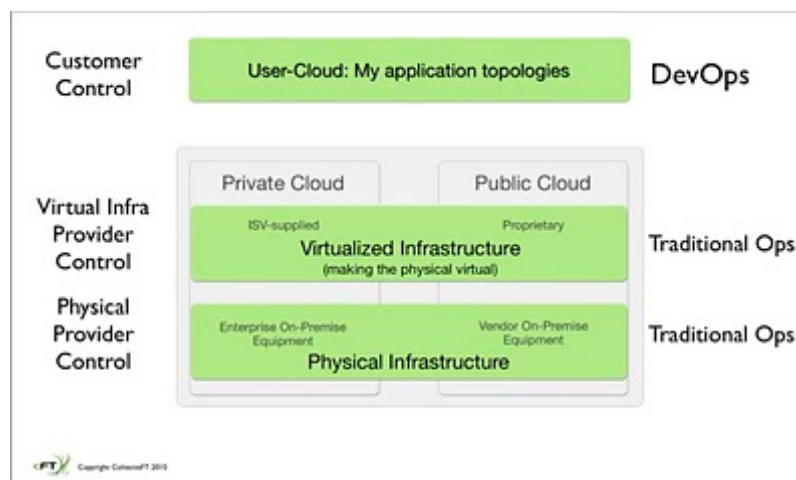
If you think about it, when you are using a public cloud like Amazon Web Services (AWS) – the team of people running the physical infrastructure have no direct interaction with your application topology; and you don't want them to.

The teams of people running your virtual infrastructure at AWS have ALMOST no interaction with your application topology; and you don't want them to. They see your topology from “below”, they see a clump of x86 VM workloads, and have no knowledge of the semantics of your application topology, and in fact shouldn't.

You, in the user-cloud, know the semantics of your application, you know how it scales, you know where it came from, where it is going, and whether it can migrate from region to region or country to country. And to date you have used either your own self-developed knowledge of EC2, private Eucalyptus, Terremark, etc., or you used Rightscale or CohesiveFT, or one of the other tooling sets that fit in the category Gartner calls “Cloud Management Platform”.

I believe this is an industry wide trend, and basically creates a new “separation of concerns” in how IT functionality is provided both by vendors and IT staffs, and how IT functionality is consumed (pictured below).

But this gives rise to the question....



Where should my virtual infrastructure and its providers fit in?

We know that a Public Cloud as my underlying physical infra provider cannot take any specific knowledgeable action on my applications, nor do I want them to. It is the same with Public Cloud companies in the role as my virtual infrastructure provider, other than a nominal number of very narrow geographic recovery actions.

What about in my Private Cloud? Today the Private Cloud (agile virtual infrastructures running in my own data centers or co-lo centers) is often mimicking existing centralized IT behaviors. In this model, a line of business (LOB) or application owner has to go “hat in hand” to centralized IT to make deployments happen.

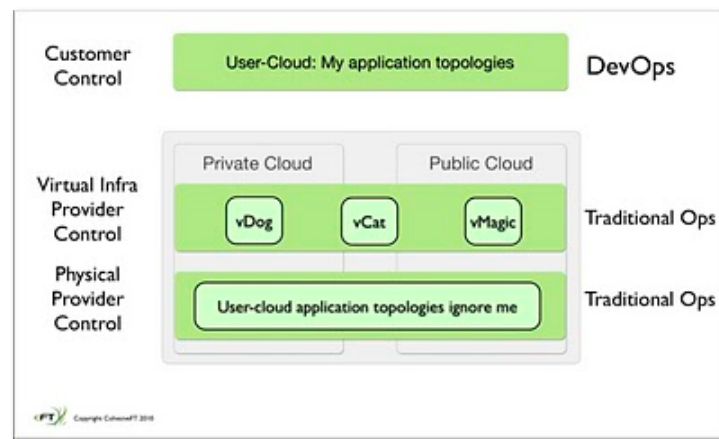
I would argue this is an evolutionary artifact which will succumb in the face of Enterprise IT successfully evolving into comprehensive virtual infra providers. Imagine if you will, the “private cloud pixies” come overnight and magically transform all of your available x86 hardware stock into one global, integrated / federated, api-driven, virtualized fabric. All of a sudden you would realize how the “frozen chunks of metal” in your physical layer (servers, storage arrays, routers, switches, etc.) had come to define the structure of your business, your control mechanisms, and your attempts at knowledgeable centralization of the application layers. As these devices become only a conveyance for virtualized compute, virtualized storage and virtualized networking – instantly the application layer accelerates to a speed and fluidity that provides poor grasp for IT staff and even less tractable control for an ISV infrastructure that perceives reality from below.

Application owners, managing from “above” know their application, its topology and are best able to deploy it to some set of described resources. These described resources (available VM slots at the simplest) probably are detailed with respect to computer power, memory availability, geography if not legal jurisdiction, storage availability, LAN speed, network edge speed, etc..

The IT operations team (which could report to centralized IT or the LOB) which knows the application best, is best positioned to bring agility and satisfaction to the application owners.

This is the next battleground for the hearts and minds of the buyers of virtual infrastructure in the enterprise. On one hand, public cloud vendors viscerally perceive the new separation of concerns. They do not want their physical infra or virtual infra staff concerned with the discrete details of any one customer's application workload.

Yet, the private cloud vendors and the incumbent virtual infrastructure providers (centralized Enterprise IT) are aligned in wanting to manage the application topology (and accordingly the Line of Business) from the middle as it were. For the IT team running the virtual infrastructure – they are avoiding “giving up” of something they fear to give up. For the virtual infra ISV, they avoid having to find and negotiate with a new buyer. They avoid having to engineer products from the top, that can integrate with their products from the middle.



This picture shows the current trend. Virtual infra providers bringing to market feature sets from the middle (as pictured vDog, vCat and vMagic) putatively to meet the needs of the user-cloud, the cloud tenants.

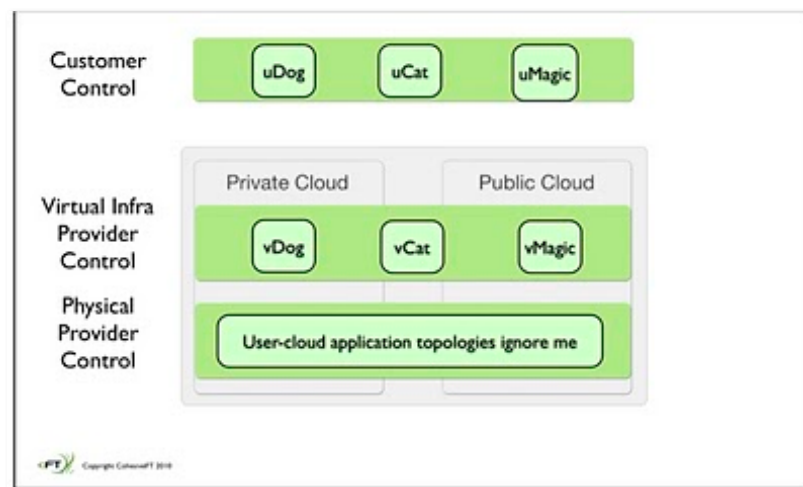
I don't like this approach. Perhaps, in a self serving sense, I am against this because it means more competition for CFT and others in the Gartner-dubbed "Cloud Management Platform space". But it is more than that. It is a violation of the new and powerful separation of concerns which is playing out quite nicely in AWS / Public Cloud space today and should have its chance in the private cloud space as well.

I am not against the virtual infra ISVs coming up with products for the user-cloud. I am happy for them to deliver a line of “u” products (user-cloud products); uDog, uCat and uMagic that use the public, published APIs of their “v” product counterparts.

What are examples of good "vDogs" or "uDog" counterparts? For one I would say Amazon's SimpleDB. It is a platform service provided from the middle, owned, operated and controlled from the middle - and can be used by application creators to create their own user-cloud applications which they can control. From the same vendor comes a good "uDog", Amazon RDS (Relational Database Service). It is provided at the user-cloud level and it is managed and controlled by the application owner. Another good "vDog" is Amazon VPC (virtual private

cloud). Again, it is owned and operated in the middle, powerful, yet comes with lessened customer control just like its database counterpart. Amazon's ecosystem provides the corresponding "uDog", CohesiveFT's VPN-Cubed; owned, operated, and controlled in the user-cloud level. As I said in my previous post, horses for courses (vHorses?). I won't point out bad vDogs and vCats here - but they are out there and I recommend thinking through their usage and their violation of the operational planes of control public and private cloud is now affording us.

When we give the cloud application owners control from above, and virtual infrastructure providers stick to managing at the workload level from below, I believe it encourages innovation in enterprise IT organization, innovation in features and functions for managing application topologies, and better interoperability and integration between cloud vendors; whether from the virtual infra or user-cloud perspective.



If we take advantage of the new separation of concerns being afforded us by the emergence of three distinct planes of operational control; physical, virtual and user, then I think we have a chance of maximizing enterprise agility and ROI as they engage in the world of public, private and hybrid cloud infrastructures.

If you are looking for me, you will find me with the rest of the CFT team, in the user-cloud.

18 September 2011

AMQP – the enchanted corner of SOA

Chris Swan - The State of Me



I first drew this chart back around 2004 for my friend Alexis Richardson. At the time I referred to it in the context of a proprietary research methodology, but I don't want trademark lawyers chasing me – hence the thesaurised title for this post.

Async	SMTP	?
Sync	HTTP	IIOP
	Unreliable	Reliable

The point was very simple – we had standards based protocols for everything except the most useful case – asynchronous and reliable. There were of course protocols that occupied that corner, but they were proprietary – fine if you're doing something within an organisation and willing to pay for it, but not so good if you're trying to get services to work across organisational boundaries.

At the time many people were trying to bend HTTP out of shape so that it could be asynchronous and reliable with efforts like WS-RM, but I felt that what the world really needed was a simple standards based protocol for message oriented middleware (MOM).

Of course my prayers were answered not soon after, with the announcement by my friend John Davies at the Web Services on Wall Street conference[1] of the impending release of AMQP, and Alexis went on to found the hugely popular and successful RabbitMQ. The brave new world was going to look like this:

Async	SMTP	AMQP
Sync	HTTP	IIOP
	Unreliable	Reliable

So why am I writing about this 7 years after first scribbling on a whiteboard? Well... it seems that some people still don't get the joke. I found myself drawing the same chart again just this week whilst in the midst of a discussion with a SaaS provider on notification mechanisms. The unpleasant surprise was that people found it useful and informative (rather than obvious and patronising).

I'm going along to PubSub Huddle later this week, where I'm sure I'll find loads of people doing cool stuff with AMQP. I just hope that I don't have to wait too much longer for this all to become mainstream. For some time I've shared the vision of John O'Hara, and others in the working group, of AMQP becoming the ubiquitous interconnect[2]. Hopefully the time for that vision to become reality is upon us now that AMQP 1.0 is out the door.

[1] I recall that John didn't have approval to talk in public about what was at the time an internal project, but he went ahead anyway.

[2] John did a great piece in ACM's Queue magazine on this (warning PDF and some scrolling required).

The original post is available on Chris Swan's personal weblog: <http://blog.thestateofme.com/2011/09/18/amqp-the-enchanted-corner-of-soa/>

Dec 30, 2011

2011 in Cloud Computing - The Empire Strikes Back!

Patrick Kerpan - VMBlog



Looking back on my 2010 predictions, I say I went 5 for 10, my colleagues insist it is 7 for 10, you be the judge. Where

I missed was actually believing that cloud-washing would start winding down (instead of ramping up) and that the basic and simple structures of the cloud market would become evident (unfortunately opacity has increased as some vendors and analysts pursue a "scare the enterprise users" approach to their business ambitions).

(As a note, I am liberal in my cloud definition, so when I say "cloud" I mean public, private, partner, hybrid, etc.. and am talking about Infrastructure as a Service. All standard disclaimers apply. While my views are obviously informed by my work at CohesiveFT, the predictions and commentary below are mine and do not reflect corporate statements from my employer.)

The "hardware support for software encryption" conversation will NOT begin. This is an about face from last year's predictions. No matter to what degree enterprises CANNOT have their data in motion in "plain text" on a 3rd party network over which they have no insight, visibility or control, the hardware side of the technology market will expend resources on security approaches that have the potential for providing customer lock-in, over solutions which would leverage existing customer skills, understanding and IT assets.

The cloud approach begins to show profound implications in how enterprises organize their IT functions. The cloud data center is becoming a three-layer architecture with three distinct planes of operational control. The bottom layer is the physical layer. Its job is to be fast, fat, and flat. It runs state of the art hardware that can be multiplexed by the middle layer (the virtual infra layer) to make the physical virtual. The virtual infra layer has a pretty focused job too; get virtualized compute and storage into play for users of the infra, and be easy to own, manage and operate for the provider of the infra. The top layer, is the "user-cloud" layer. It is from this perspective that individuals, IT operations and lines of business deploy business capabilities running on top of the virtual infrastructure layer. These three layers have distinct interfaces, and a clean separation of organizational concerns. This allows enterprise IT a dramatically new framework for managing expertise, staff, service partners and software and hardware vendors.

Savvy enterprise buyers will tell their hypervisor vendors to "shut up and drive". Like a nosy taxi cab driver or limo driver, hypervisor vendors are involving themselves in additional feature and function areas in ways which disrupt and deform the powerful new separation of concerns in the enterprise cloud datacenter. As hypervisor price commoditization has occurred, the providers of hypervisors are looking for ways of returning that "not quite so thin anymore layer of software" to a position of power. In the interim they have focused on making pools of hypervisors easier to own, manage and operate. This is a good thing! They are awakening to the critical need for federation (making pools of data centers of hypervisors easier to own, manage and operate). This is good too! However, they are also trying to make the tightly controlled hypervisor a nexus of control over their customers. With traditional server software - customers have complete control of what they do and don't run in the server OS. Enterprises can make it as lean or as fat as works for their organization. Hypervisor vendors control their stripped down OS that runs on customer hardware - and they make decisions about what can run in it. And now hypervisor vendors are telling customers they (the hypervisor vendor) will make all sorts of single tenant, single choice, unilateral decisions about what you need in that layer of your infrastructure. Savvy buyers - beware.

"Cost Zombies" start to be a concern. Many a virtual machine "sprawl management" startup has come and gone over the last few years, and as a rule I still think we don't have enough virtual machines out there. But it is possible now to go "ouch, \$500 this month on some cloud virtual machines I had forgotten about". As enterprises become more facile with agile infrastructure expect them to demand time and dollar budgets to be part of the cloud server launch parameters. This could be provided and enforced by either your virtual infra provider or your user-cloud providers. As invoicing and purchase orders join the credit card in the public cloud offerings, it will be interesting to see the emergence of these strategies.

"Federation" becomes a big part of the enterprise virtual infrastructure conversation.

Providers of virtual infrastructure such as IBM, VMware, Citrix, and Eucalyptus obviously have intellectually understood how many data centers a large enterprise has. This year begins the emotional and visceral understanding of what this means and how they need to appear as continuous fabrics as opposed to islands of hypervisor and virtual machine isolation.

Microsoft has an "aha" moment around Amazon-style services via the Azure "VM role".

Azure has fought a retreat since its inception, resisting being an infrastructure as a service offering that can run X86 workloads defined, controlled, manufactured and uploaded by

customers. It began as a "language runtime cloud" offering a set of platform services which could be consumed by your .NET byte code and executed "in the cloud". Virtual machines had no role. Unfortunately, and somewhat obviously, enterprises couldn't even begin to imagine the attack vectors of such an approach much less envision the defense, governance and audit approaches necessary for it. Then came the "well, it is a virtual machine, and you can pay-per-hour for a little, middle or big one, upload your .NET code, control some parts of it, but no you cannot fully control it as a virtual machine" offering. There still were not a lot of takers for that approach. Now they have announced the Azure "VM Role" which allows customers to upload their pre-manufactured Windows Server Virtual Machines and run them. Customer's have wanted and waited for this - and with the investment Microsoft has made in its Azure datacenters - customers will use this. Customer acceptance of this will light up the bulb over Microsoft's head and we should see some good come of it.

Another public cloud offering will reach critical mass and exit 2011 poised to be "the enterprise public cloud" in 2012, and Amazon's first serious competitor. There are only a handful of possible participants from the enterprise computing space, a couple from the Internet consumer space, and some dark horse telcos around the world. I think one of them will put it all together sufficiently to go toe-to-toe with Amazon, but with an enterprise focus. At least one of these players (my money is on IBM), will get through the business model complexities of cannibalization and sales compensation in 2011 and be strongly positioned as the "enterprise grade" public cloud. If not, then Amazon will have a free pass to become the dominating force of both the Web 2.0 and the Enterprise cloud markets.

We will begin to speculate about the first Cloud IPO. The name most mentioned in this topic will be RightScale. Fair disclosure, my company, CohesiveFT, partners with RightScale in a number of accounts, competes in other accounts, and are a non-intersecting set in others given differences in feature set and target customers. That said, from a financing and valuation metric (if one believes the rumor mill), they appear to be the leader of all the private companies/startups in the entire cloud stack. As the recession recedes, and the cloud market matures, they have a real shot to be the first big investment success of the cloud era.

Originally published in the VMblog Virtualization and Cloud Prediction Series. For the original post, visit (<http://vmblog.com/archive/2010/12/30/cohesiveft-2011-in-cloud-computing-the-empire-strikes-back.aspx>) A rights reserved by the author.



The First Drops: Cloud in Production: 2012 - 2014

Exclusive Viewpoint: Cloud Still Has a Long Way to Come

Chad Lawler

Director, Consulting Services, Cloud Computing at Hitachi Consulting



It's very interesting to look ahead and try to anticipate where the puck is going. The truth is that what Werner Vogels has stated about where we are in the "maturity curve" is absolutely true – that even with the progress we've made, we are very early in the cloud evolution. We want instant gratification with a solution that fits all of our needs, but we're not there yet. We're certainly not at the point where true utility computing is a reality.

Here is an example: most cloud vendors say they offer pay-per-use functionality, when really they're offering pay-per-allocation. You allocate a small, medium, large, or extra-large compute size, server memory, and a certain amount of storage. You pay for what you allocate, not what you use. When you allocate 10TB of storage on AWS S3 and do not use all of it, you still pay for the 10TB allocated. In order to have scalable storage you still have to pay to allocate it.

The idea of utility computing is still very relevant, interesting, and appealing to businesses, particularly when you get into the potential for dynamic, hybrid cloud. With hybrid cloud, you can get to daydreaming about moving workloads on the fly to the cheapest or fastest cloud.

The truth is, that just isn't a reality today. It's coming, however, as you can see with recent releases by leading cloud providers that enable basic hybrid functionality. However, the standards are not there yet and clouds are proprietary. It is very complex to move a traditional application from your data center to a public cloud and back without detailed knowledge of the application. Plus, there are still many integration, performance, data consistency, and security concerns, and not every application is a fit for the cloud, at least today.

Cloud will continue to evolve toward a hybrid, utility-compute cloud model, but it is going to take a long time; I would guess 7 to 10 or so. Entrenched technology vendors don't want to erode their existing revenues by creating competitive cloud offerings, so they resist change and cloudwash as a result. This is very confusing for business and customers. Many other technology vendors are digging in their heels and developing very compelling private cloud offerings to retain customer control. All of these complex trends are creating a "push and pull" scenario between the true utility computing we want in the long-term and the status quo of where we are today.

Wednesday, September 5, 2012

After VMWorld, who's talking about the Application layer?

Margaret Walker - CohesiveFT Blog

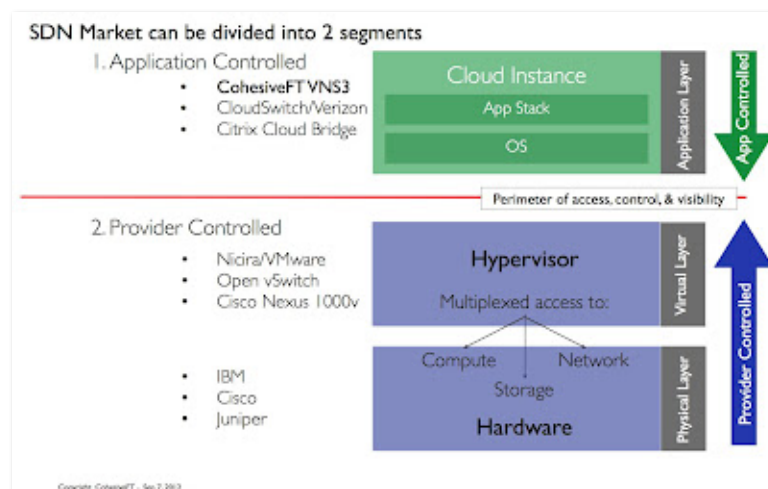


Another VMworld has come and gone. Now the blogrolls are filled with chatter about all the new features of storage, open source software foundations, and jockeying among cloud providers.

Don't get us wrong, the Cloud Service Provider (CSP) market with it's hypervisor and hardware innovations is extremely important. Without the provider market continually making the cloud/virtual layers more fast, fat, and flat there would be no users running applications in the cloud. No customers means we wouldn't be able to build our business in these cloud ecosystems.

We support the awesome innovations and discussion out there and are excited for the future. But we think it's time to expand the conversation to include the people and the workloads they are running in the cloud today.

The Application Layer



Who is talking about the application layer? We want to talk more about the user-controlled, "cloud service user" (read: application) layer and would love some company. As an added benefit, the cloud user would definitely appreciate the application layer receiving a little more attention. Who is in control of choosing the images (OS and kernels), connectivity, and topology context for a given cloud deployment? It's definitely not the cloud provider.

Wednesday, September 26, 2012

OpenFlow is Software Defined Networking, Software Defined Networking is not just OpenFlow

Patrick Kerpan - CohesiveFT Blog

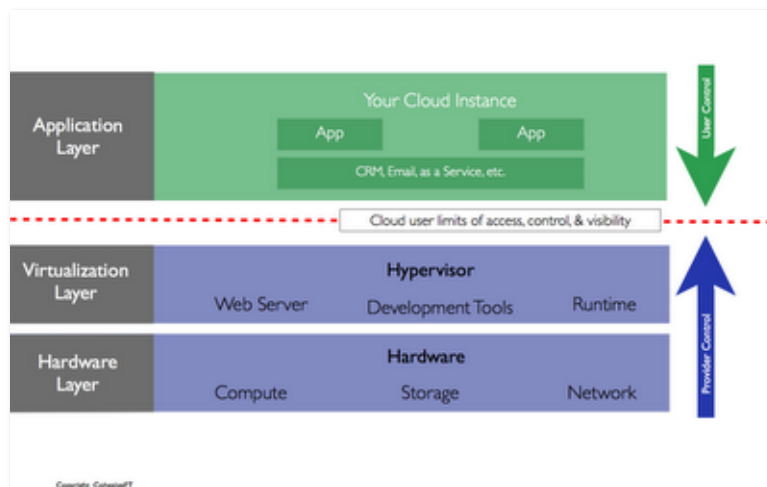


Software Defined Networking and the Long Road to SDN's Overnight Popularity

With its recent acquisition of Nicira, VMware has "done it again"; they have taken a space and made it an "overnight success". The first time they did this was when EMC floated the IPO of VMware in 2007. A company that sowed the seeds and did the spade work for x86 virtualization over the course of a back-breaking 9 years, made itself and the topic of computer virtualization, an "overnight success" to those who hadn't been watching, learning and participating over the course of those years.

SDN solves the architecture and configuration issues of today

HP's Bethany Mayer spoke about SDN's solutions for current hurdles. Mayer said, "the network today is rigid, architected for a single tenant. It has to be architected for more than one user. You have to go box-to-box to change the configuration. And 70% of errors occur in the manual entry of CLI (commands)." The solution? SDN's ability to connect and separate points of control for the application and hardware.



We use the above image to illustrate the control issues by layer. Like Mayer mentioned in her speech, the key benefit of SDN is the application-layer control. The access, networking and security features can be controlled by application users because the needs and concerns of the

cloud service provider are distinctly different than the needs and concerns of the cloud service user (the application topology deployed to the cloud and its owner).

The 7 Properties in the Evolution of SDN

The founders of Nicira, now part of VMware, wrote one of the original pieces surrounding the SDN concept, the Seven Properties of Network Virtualization. They set quite a high hurdle by declaring, "network virtualization, if done correctly, should be able to run any workload that is compatible with existing networks, over any type of hardware at any location. The following list of seven properties must be in place to gain these benefits. Without them, it is not possible to unlock the true potential of cloud."

The 7 properties, in summary, that news-creating Nicira founders, Open Networking Foundation folks and CohesiveFT agree on are:

- Independence from network hardware
- Faithful reproduction of the physical network service model
- Follow operational model of compute virtualization
- Compatible with any hypervisor platform
- Secure isolation between virtual networks, the physical network, and the control plane
- Cloud performance and scale
- Programmatic network provisioning and control

"Big Tent" Thinking and the The Future of SDN

As techniques like OpenFlow gain a foothold in the physical and virtual infrastructure layers and the multi-tenancy models mature, then OpenFlow conversations should be taking place between the provider-controlled layer and the application-controlled layer. CohesiveFT holds that we should continue and expand on the 7 properties and use "big tent" thinking to capture the momentum of SDN and virtualization.

Wednesday, October 17, 2012

The future of your data: Cloud, Fog, and Flood

Patrick Kerpan - CohesiveFT Blog



How the cloud will extend to the ground to cache everywhere for everything, and how data will change us all.

The cloud gives all business, regardless of industry, the flexibility to economize data, hardware and software through third-party infrastructure. Cloud service models are the ways you see the cloud in your daily life and are the acronyms you've probably heard of: SaaS, PaaS and IaaS. If you want more background into the basics of IaaS, PaaS, and SaaS check out this great article from VentureBeat.com.

In business, access and connection to data - via the cloud - are changing where and how we work and speeding up product development time-to-market. Back in the day, (traditional) architects created a CAD model of a building, burned the huge file to a CD or disc and mailed it to engineers, who then edited and returned it before any brick could be set. Now firms can collaborate to share updates back and forth via hosted FTP and dramatically shorten the time from design stage to a built structure.

People ask us frequently some version of the same question, Where do you see cloud computing in 10 years? Our short answer is the slightly cryptic, "Cloud, fog and flood: cloud to the ground, cache everywhere, my data everywhere always." Here's the deeper explanation of what that means to us:

Cloud to the ground: Data behind every wonderful app in your daily life

At the very basic level, cloud masks the technologically complex work that goes into providing the features, functions and access to cloud-based applications. When you upload files to Dropbox.com and later retrieve them from another computer, you don't quite know how it happened but it makes sharing data much easier. Beyond technology, though, the cloud is already starting to impact much more than the IT department.

Accessibility to your applications such as web-based email, Dropbox, Salesforce and other wonderful cloud-centric apps (SaaS in this case) lets you do work at home, from the coffee shop or collaborate around the world. No more software discs, network ethernet tethers, or data server closets. Case in point, the cloud takes away the need for Netflix to run and manage its own building full of hardware data centers and frees IT departments to focus on core computing and data tasks.

Data fog: Cache everything for better information

Most prevalent in web applications, data caches store previous information on web browsing history or websites you go to most often. As in, the “remember me” options on ever login, your fantasy team lineups from 2 years ago, or Google Chrome’s Most Visited pages. The data is close at hand to streamline user experience, but what about bigger data needs? What about the 1M gigabytes a film studio needs to keep on hand when creating a 3-D movie? That’s a lot of data cache.

In your office the fog of data - no matter if your business manufactures widgets, publishes online cookbooks or sells alternative power - requires the cloud computing capabilities to access, store and manage the additional data on-demand. Updating and syncing data will create giant caches, which ratchets up the need for data storage and cloud computing oomph.

Cloud flood: The rising tides of data beyond the IT department

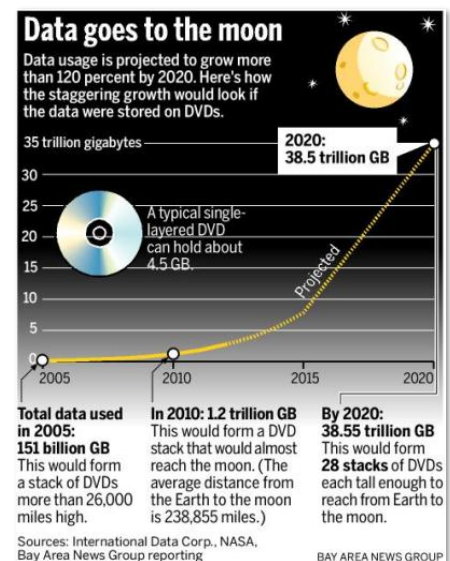
Data is already creeping into daily life, but will soon flood most users because we are demanding access to “all my data everywhere, all the time.” Mind-blowing stats include:

- More video was uploaded to YouTube in the last 2 months than if ABC, NBC, and CBS had all been airing new content 24/7 since 1948 (more on internet stats).
- Since 9/11, the U.S. Military’s remotely piloted drones and other surveillance technologies have gathered 1,600% more intelligence data.
- The Economist estimates mankind created 150 exabytes (billion gigabytes) of data in 2005 but in 2010 it will create 1,200 exabytes.

In the enterprise, the access to data also creates new customer insights and tracking like never before. Marketing departments* use data to influence sales from online shopping history, Netflix’s suggestions, and Target’s (overly) insightful mailers. 25% of respondents to a Gartner survey have invested in big data projects. Another 31% plan to do so in the next 2 years. But what does the coming flood mean for your organization?

Prepare for the flood: Virtualization and Cloud Migrations

According to Forrester Research’, 85% of North American enterprises have already adopted x86 server virtualization or are planning to expand their implementation in the next 12



months. In fact, only 4% of respondents are not planning to implement x86 server virtualization in the future. Only 48% are virtualized today.

Cloud computing is so alluring because it economizes the computing resources and data centers. Rather than own and manage (and pay utilities for) organizations' data center or hardware needs, companies can use the cloud model to scale computing needs to business needs. In a NY Times article, David Cappuccio, managing VP and chief of research at Gartner, said his own recent survey of a large sample of enterprise-run data centers found that typical utilizations ran from 7 - 12%. With multi-tenant computing, enterprises don't have to figure out how to optimize idle data servers, but rather "rent as you go" from a cloud provider.

Cloud to the ground and your business

The infusion of data, transport and computation into the fiber of all business means working with the cloud in all its forms will be critical. Cloud to the ground has the potential to be so ubiquitous - it will both enable and drive the need to use the flood of data to make decisions, automate workflows, and build new business. Of course "cloud" doesn't replace all of the normal business operations, without which such infrastructure is meaningless. But it does mean new forms of business radar, running lights, and sonar are needed to thrive as opposed to merely survive.

October 18, 2012

Security and Control in the Cloud: Three migration rules to break

Patrick Kerpan - Cloud Computing Journal



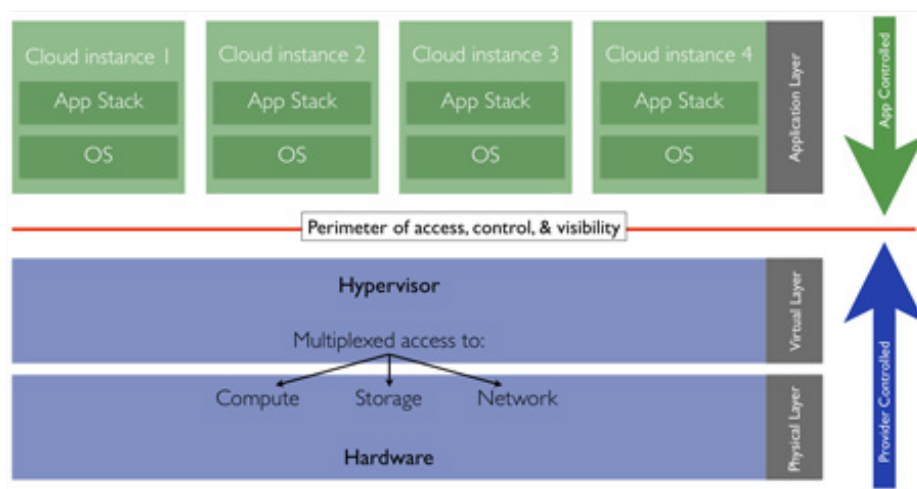
Cloud computing is so alluring. The public cloud economizes infrastructure resources and creates a scalable, on-demand source for compute capacity. Additionally, the cloud can be a strategic asset for enterprises that know

how to migrate, integrate and govern deployments securely.

Apple co-founder, Steve Wozniak recently said, "A lot of people feel 'Oh, everything is really on my computer,' but I say the more we transfer everything onto the web, onto the cloud, the less we're going to have control over it." In fact, over 70% of IT professionals worry about security according to an IDG Enterprise Cloud Computing Study.

Boiled down, security, access and connectivity are really issues of control. As any prudent cloud user, the application has its own unique security features, such as disk encryption and port filtering. But do these layers of security features overlap or conflict? What happens to ownership after migration? Do solutions really have to be architected before and after deployment?

Take an application-focused approach to security from the beginning. The application-controlled, application-owned security layers will ease the decision to deploy, test, and develop in the cloud and save on IT training and time along the road.



Control of Security: Who Has It?

Part of the "magic" cloud providers and vendors supply is wrapped up in layers of ownership and control in the form of firewalls, isolation, and the cloud edge. Most enterprise application owners hope that these layers will cover the possible gaps in security after migration.

Unfortunately most enterprises need security controls they can attest to and providers ultimately own and control these security features.

Unfortunately the needs and concerns of the cloud service provider are distinctly different than the needs and concerns of the enterprise cloud service user (the application topology deployed to the cloud and its owner). Security loopholes can exist because there are gaps between the areas users and providers control and own. The known boundary between what the cloud user can control and view and what the cloud provider can view and control is the root source of enterprise executives' concerns with public cloud.

The provider-owned, provider-controlled features (as in the cloud edge, cloud isolation), the provider-owned, user-controlled features (or the multi-tenant API controlled router / hypervisor), and the app-owner, app-controlled features (OS port filtering and disk encryption) can be configured in an overlay network to give the user the ultimate control of security.

Application-to-cloud migration and software defined networking (SDN) capabilities out there offer additional, overlapping layers of control and security that span the spheres of the traditional cloud layers.

In order for cloud projects to succeed, IT executives need methods and tools they can attest to and can pass audit. Understanding the perimeter of access, control, and visibility between the application layer and the cloud provider layers is the first step to a less painful cloud migration. With this knowledge enterprises can then design a migration process that fits their use-case to deploy application topologies to the public cloud in a secure and controlled fashion.



Three Migration Rules We Recommend Breaking

Today's migration "rules" create more hurdles than solutions. Rapid industry changes, lack of standard security approaches, and the confusion on the proper steps to cloud deployment cause enterprises to overlook the issues of application-level control.

In fact, application-centric concerns are not even being addressed. Popular migration advice urges enterprises to tackle huge hurdles before and during migration, including deploying all at once, re-architecting before migration, and postponing the cost benefits of using the cloud.

Break the following three migration rules and it is possible to renovate more efficiently, capitalize on the cloud's economies of scale, and quickly, easily, and securely control enterprise networks and applications in the cloud.

Rule 1: Deploy all at once or not at all

Just as lemmings became extinct by all jumping in head first, most enterprises require time to analyze and adjust to new technologies before committing serious time and effort. Employees, customers, and shareholders would not be happy if companies jumped into new technologies without first proving value. Thankfully, enough enterprises, organizations and governments have already seized the benefits of the cloud's flexibility, cost savings, and connectivity.

Now, the challenge for IT professionals is to find the cloud architecture and provider(s) that fit their enterprise's needs and avoid having to reinvent the cloud to do so. With proven solutions in the market, enterprises can skip the bare metal to virtual to test cloud development life cycle. Simply deploy directly to any cloud environment, develop, test, then release to speed the time to market.

Rule 2: Re-architect before migration

Most providers and brokers want enterprises to spend time and effort to re-build IT systems and as a result re-learn / re-train before migration. Advice articles list migration steps of parsing applications, virtualizing, re-architecting and then migrating. Cloud pundits advise IT professionals to be wary of all cloud security and take valuable time to renovate before migrating - which will slow down the process and postpone or even wipe out the financial benefits of the cloud.

The traditional datacenter has too much knowledge flowing in a vertical direction from application to infrastructure and infrastructure to application. Migrating to the cloud before the renovate, design, or innovate steps can cut down on the upfront hassle by removing the burdens of re-architecting and re-learning skills before migration. Saving time, IT resources,

and forgoing the arduous re-training speeds up the process for migrating to the cloud and ultimately how the organization capitalizes on the cloud's flexibility.

Rule 3: Pay upfront for design and renovation costs

Why stop with the cloud's physical economies of scale when there are potential savings on the costs of IT overhead? The same time and effort put into saving "design economies of scale" can be used to save major overhead costs too. A single migration, rather than the process of backup, re-architecture, and then migration is more cost-effective. Why wait for cost savings until after migration when there is an option to realize faster deployment and speed to market?

The added customization and control needed to migrate in a logical set of steps puts the control and security solidly back into the application layer.

Enterprises will likely face a long, slow migration to the cloud but, with the tools to capture the efficiency of migrating through logical steps before designing, the process can be significantly less painful. The application-controlled, application-owned security layers will ease the decision to deploy, test, and develop in the cloud and save on IT training and time along the road.

Conventional wisdom is missing the application layer importance of security and control in the cloud. So only one migration question remains - why take the stairs when you can take the elevator?

Copyright © 2012 SYS-CON Media, Inc. — All Rights Reserved. For the original, visit: <http://cloudcomputing.sys-con.com/node/2402550>. Syndicated stories and blog feeds, all rights reserved by the author.

Wednesday, November 28, 2012

Why you should Waste Infrastructure: DevTest is Dead

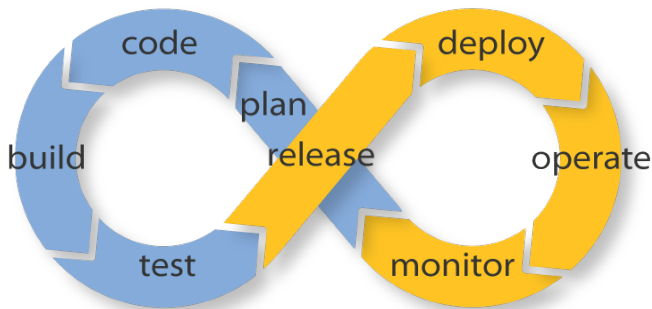
Margaret Walker - CohesiveFT Blog



Enterprises, organizations and governments have seized the benefits of the cloud's flexibility, cost savings, and connectivity. Yet the traditional approach to dev/ops is decelerating cloud economies of scale. The known physical economies of scale focus on saving enterprises and customers the labor and maintenance costs of managing the data center's physical footprint. The next

level of savings is focused on the production footprint - the cloud as the new design center.

Design Economies of Scale



Endless Possibilities: DevOps can create an infinite loop of release and feedback for all your code and deployment targets.

Why stop with the physical economies of scale when you can save on the costs of design overhead? Save time, IT resources and capture the design economies of scale: migrate, then renovate/innovate in an environment that is both agile and secure. With the cloud-based design center, the concept of saving precious infrastructure has been replaced with a need to

save the development team's precious time.

Most developers cannot possibly predict the needs, capabilities and budgets of their customers when beginning a new project. Currently, developers are building out services for the cloud but then worrying about optimizing their app for a certain cloud, a certain criteria or a different optimization that comes at a lower price point. What is a developer to do when looking down the barrel of rapidly evolving cloud targets, sensitive budgets and enterprise IT administration?

Waste Infrastructure in the Cloud Design Center

There is no need to lose sleep over how to correctly forecast the architecture and infrastructure before building. The new cloud design center offer a new way of developing, building and testing new instances and resources as needed.

Developers can start with production and use the flexibility of the cloud to build and tear down production footprints to use more fast, powerful resources to keep design stages innovative and geared toward the target production. Quickly move past the old steps of

develop, test, migrate, deploy then re-architect. Cloud design centers also allow multiple parallel streams of work on identical yet distinct copies of application topology.

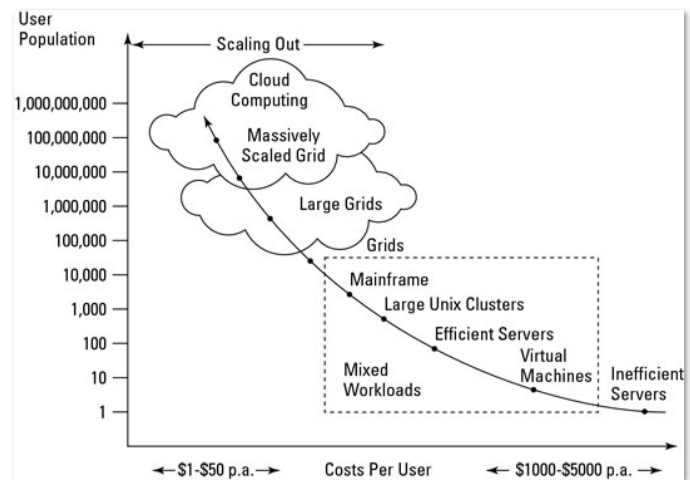
An international trust used this model to begin with a realistic development and infrastructure in mind. The greenfield project began immediately as both a development and testing production stage, capitalizing on the plentiful infrastructure to define a faster route to a realistic target environment.

Get in the Cloud Express Lane

Going through the traditional dev / test steps before migrating to the cloud is similar to

sitting in deadlock traffic while cloud designers are zooming by in the the express lanes. Skip the worries and tie production and development closer together with the automation and testing in the cloud environments, not to mention the time you can save from working in a safe environment rather than risking any “dark ops” security issues. Development needs maximum scalability, redundancy and availability, the cornerstones of cloud offerings, now with agile guidelines to open up to innovation rather than narrowing by avoiding risk.

Will you be stuck in traffic, or can the cloud offer you a ticket to the express lane of development?



Tuesday, June 11, 2013

Software Defined Networking and Network Function Virtualization

Chris Swan - Wired Innovation Insights



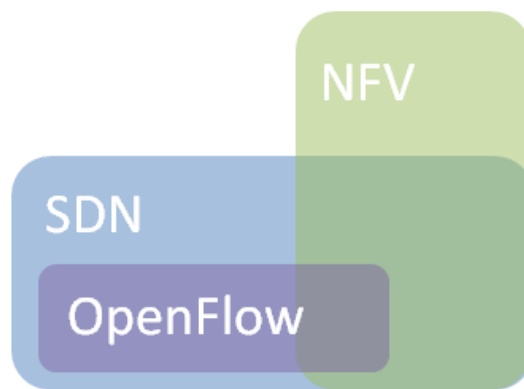
Software Defined Networking (SDN) has been a hot topic for a little while now, and remains of great interest as people try to figure out approaches to a Software Defined Data Center

(SDDC).

Network Function Virtualization (NFV) is a less familiar term to many, but is a good label for some important developments where networking meets virtualization and cloud computing. SDN is where an application programming interface (API) is used for configuration and management of a network. NFV is where commodity hardware is used to run networking functions in software.

The Overlap

Since both SDN and NFV relate to networking and software there's quite an overlap between the two things, especially when it comes to talking about implementations:



Since NFV is implementing the network in software it would be somewhat ridiculous to have a good old command line interface (CLI) as the only means for configuration, but that doesn't mean it's not possible. Most NFV implementations come with an API (and quite often a web interface built on top of that API), but that's in no way mandatory.

Similarly Software Defined Networking doesn't have to mean software implemented networking. For sure there are a bunch of SDN implementations that have software moving packets around, but there are also plenty of SDN implementations that are hardware based.

Standards

The standard that everybody seems to be talking about in the SDN space is OpenFlow (to the extent that for some it seems like SDN==OpenFlow). OpenFlow has become by far the most common 'south bound' interface - between an SDN controller and a network forwarding plane. Standards for NFV are yet to emerge, though there is some work going on at ETSI. To a certain extent there's less of a need for standards with NFV as it's a classic case of same wine, different bottles. The network functions are usually the same as they've ever been, and the protocols used are too. For large scale operators (like telcos) there might be some desire to standardise aspects of the hardware and any virtualization layers used and associated management/deployment tools, but as we've seen more generally in the world of virtualization and cloud platforms (lack of) portability between implementations hasn't been much of a barrier to adoption.

What's it Good For?

The advantages of SDN have been covered at length elsewhere. For me the key point is moving away from the interface to the network being a trouble ticket to the network operations team, which is a key prerequisite to data center automation at scale.

The advantages of NFV depend more upon perspective. For large network operators it's about moving from costly proprietary models to commodity implementations, which also has impact upon scale and agility (it's much easier to source, provision and deploy onto a bunch of regular servers or VMs).

For end users NFV is about the democratisation of networking. If it's now possible to deploy a router, firewall, VPN concentrator or protocol redistributor onto a cloud where it costs a couple of cents an hour then all kinds of things can be done that were previously too difficult or costly. SDN is making networks easier to configure and manage, whilst NFV is making networks easier to deploy and scale. The two fit together quite naturally, as the key theme is software. This is especially true in the cloud, where you can't ask a service provider to install a lump of hardware for you.

Originally published by Wired Innovation Insights. For the original post, visit <http://insights.wired.com/profiles/blogs/software-defined-networking-and-network-function-virtualization%23axzz2hjonlUER>. All rights reserved by the author.

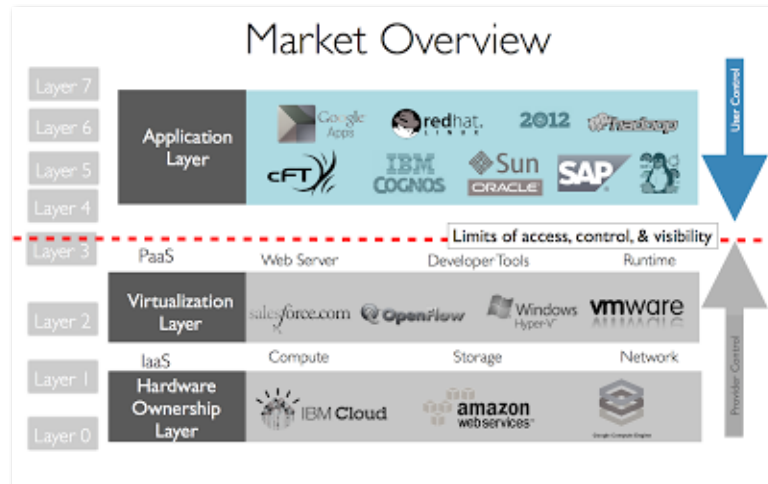
Thursday, July 25, 2013

Cloud markets and networking layers

Patrick Kerpan - CohesiveFT Blog



Let's look at the software defined networking space. It's one of the buzzwords in the market you're hearing more about, along with network virtualization, network function virtualization (NFV) and things like Overlay versus OpenFlow.



You can think about the SDN market and the feature and functions set as existing above and below this line of access and control.

Or you can say the boundary is defined by service provider infrastructure. You, the customer and your applications run on that infrastructure that the service provider offers.

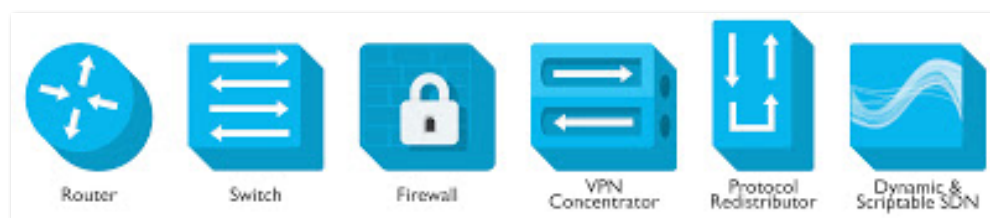
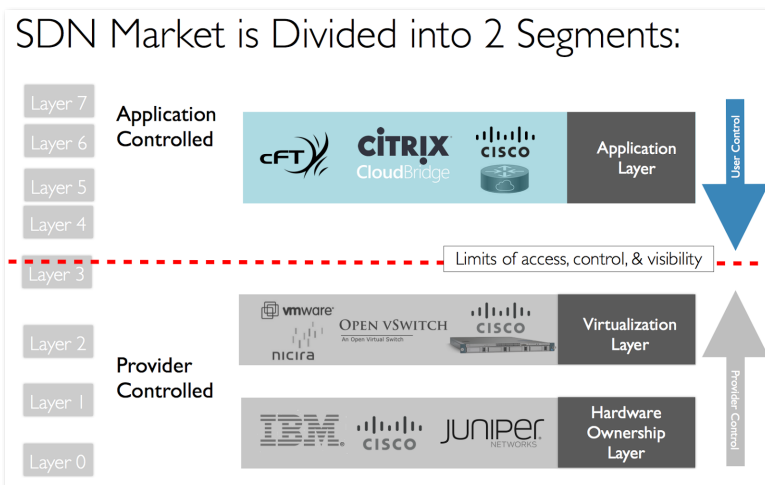
So when you look at the service provider layer, the big noise is driven by the energy around the emerging OpenFlow standards and the Nicira acquisition by VMware.

And what that has focused the market right now and in the immediate future not on the customers' ability to control and deliver applications, but has been about helping service provider deliver a better cloud. Service Provider SDN is about "help me deliver a better cloud. Help me run my data centers as a global service provision business." All those things in the infrastructure are good for customers, but not directly. There are really only indirect benefits in Service Provider SDN for most companies.

There's an incredible amount of innovation going on below the line, and service providers are growing with better service and federation. This growth is all part of the compelling reasons why cloud infrastructure is growing and growing and assuming such a dominant part of our landscape.

When we go above the line into the world of you, the customer, this is where your applications reside. You want to deliver your applications to the cloud. You want to be in control of those applications. So you tend to have solutions like ours, VNS3.

VNS3 is an instance-based solution. You run VNS3 instances as part of your overall cloud topology. You can run these as virtual routers, switches, firewalls, protocol re-distributors. They become a part of your application topology similar to how your networking devices in your data center work with your network at an individual row or rack layer.



So, you're moving a network to the cloud with your application, not necessarily the network as a whole.

So in terms of the SDN topic one thing to note - we've been doing this since October 2007. We have an extraordinary depth of insight and customer insight with our customer experience in networking and the cloud. We've driven about 60 million device hours with some of these use cases you're about to see. And about 90% of our customers are running production loads in the cloud, where their ability scale their business is based on our ability to deliver to them.

Tuesday, August 13, 2013

Our Market Landscape and an Awesome Prezi

Margaret Walker - CohesiveFT Blog



With help
from +Zane
Groshelle and
the awesome
folks at Prezi,

we present our view of the
cloud market landscape:

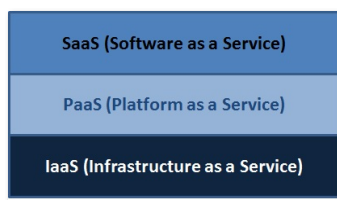
View the Prezi here, too.

Reasoning, if you will

We've tried our hand at
making complex market /

industry graphs, but worried the diagrams were too complicated, even for us. To better
quantify and explain the cloud market, we started with 3 market assumptions:

Market concept 1: the cloud stack



The basic layers of cloud computing that split the cloud market into Software as a Service, Platform as a Service, and Infrastructure as a Service.

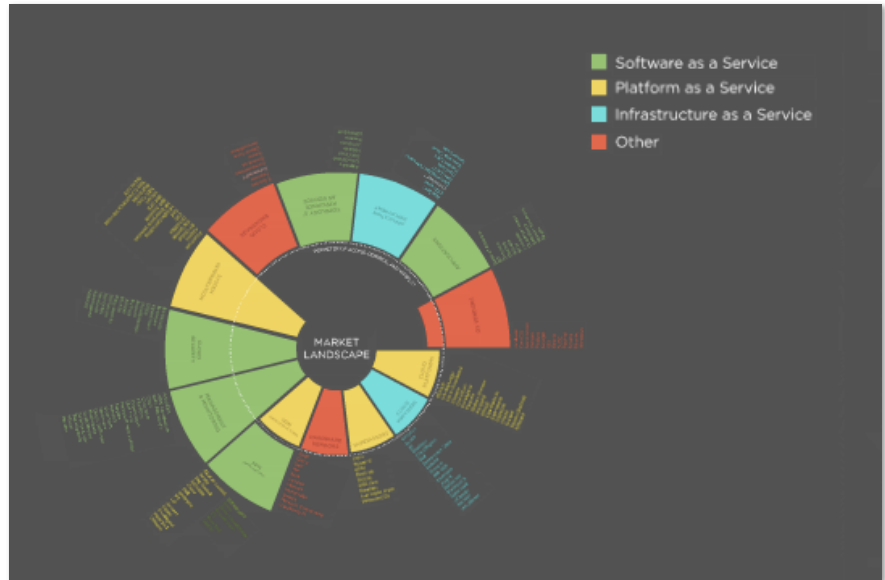
Market concept 2: networking & layers of control

References to the good old 7 layer OSI networking
dip to flesh out the cloud stack into more specific cloud offerings such as
hypervisors vs. OS vendors and cloud service brokers.

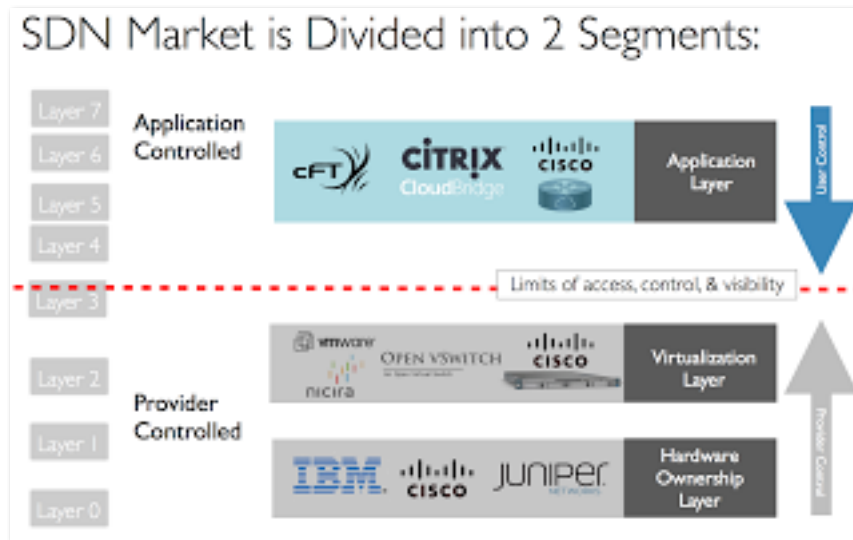
Market concept 3: the divide

Now, how to divide market segments and vendors? We spent a while debating
where each market functions relative to our concept of the perimeter of access,
control and visibility.

About that perimeter...



Our internal discussions centered around the perimeter of access, control and visibility. We made some executive decisions for how vendors operate and market. We also drew insights from Gartner analyst +Thomas Bittman, who suggested in 2010 that a key feature of cloud computing is the “boundary between the customer and the provider.”



Location, Location: SDN segments

In past blog posts, we’ve written about how Software-defined networking (SDN) can fall into two categories: OpenFlow based or and Overlay based SDN, depending on the layer focus of each.

We needed serious help organizing our market

landscape. That's where the talented folks at Prezi came in. Does our market landscape jive with your industry view?



The Future - Floods of Cloud Uses: 2014 onward

Dec 6, 2013

2014 Predictions: Amazon will lead, but the gap is closing

The CohesiveFT Team - VMblog



Amazon Web Services (AWS), the public cloud Amazon has built up since 2006 has continued to dominate the market.

Both technology features available and user base lead the competition. Gartner analyst Doug Toombs writes that AWS have over "5 times the compute capacity in use than the combined total of the next 14 largest providers (based on Gartner's estimated market share statistics)." The cloud provider's wildly popular re:Invent conference, only started in 2012, saw attendance double in 2013.

But what is the future of AWS? Will they be able to hold on to the lead for another year or six? CohesiveFT team members weigh in on the aspects on AWS that are keeping the cloud provider up and to the right, as well as other providers who can compete in 2014.

AWS will lead with innovation, but GCE and HP will follow.

Amazon will continue to innovate ahead of the rest of the market. Clearly the succession of product launches announced at re:Invent at the end of 2013 show that Amazon isn't standing still and waiting for the field to catch up.

CTO Chris Swan writes that newer clouds like Google's Compute Engine (GCE) will use architecture advantages to get an edge on AWS. Fresher architectures allow lessons to be learned from the (few) mistakes that Amazon has made. GCE's live migration is a good example - it doesn't just benefit from a more flexible hypervisor, but also takes advantage of networks that can span availability zones and regions.

Ryan Koop, Director of Marketing and Products agrees and adds that HP is also positioned to become a credible AWS challenger. HP is focusing on both public and private cloud through its use of OpenStack. Openstack has allowed them to get a head start on the public cloud offering (HPCS 13.5 based on Grizzly offers 80% of AWS VPC functionality today), gain insights and efficiencies on the public side they can take to the private, and use open standards and common infrastructure to offer hybrid cloud capabilities where other providers have failed. HP has also signaled through announcements and actions that it will build, support, and foster it's cloud partner ecosystem rather than build supplementary services themselves (see Akamai for CDN, Stacko for PaaS, Gigaspaces for SaaS, etc.). This recipe suggests success

in 2014 by allowing HP to control the customer relationship / experience, offer a high margin service and focus on its core competencies.

Rogue IT will be reigned back in.

The Abiquo team (one of CohesiveFT's partners) predicts an ebb in user-provisioned Amazon accounts and AWS independent of the IT department. Rogue teams "had a detrimental effect on the IT organisation, not to mention the unforeseen costs involved," writes Ian Finlay, VP of Products at Abiquo. "As a result we have seen a demand from the IT department to gain back control through the adoption of management platforms to gain visibility around what resources to acquire and where. IT departments started to direct attention to spending in cloud." Look for more top-down decision on cloud provider spending in 2014.

Similarly, Ryan Koop thinks the C-level will steer purchasing decisions toward the familiar brands. HP, Microsoft, and now even Google, are incredibly familiar and trusted brands in the enterprise world. Some enterprises will continue to struggle with buying IaaS from an online bookseller regardless of what their hip dev / ops team advises.

And lastly, for pricing factors in the cloud provider market, CohesiveFT team members see different outcomes.

CTO Chris Swan thinks a full on cloud price war will not break out. He writes "elasticity of supply remains way too low, and even the largest providers aren't placed to cope with the demand that lower pricing would bring. AWS will continue to haircut prices (at a lower rate than their cost of supply), and their competitors will follow. Google have started something here with its 10% cut against comparable Amazon offerings. That might be enough to sway projects that are starting from scratch, but a larger delta will be needed to make migration between services a serious consideration (and lead to counterattack price cuts)".

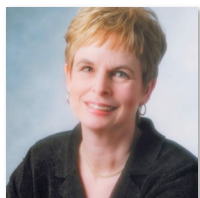
On the other side, COO Dwight Koop thinks AWS will not increase prices to match the competition. Yes its true, AWS prices are just too good to be true. In spite of the massive marketing campaign by the incumbents in the technology sector justifying their higher prices, Amazon will be too stubborn to increase its prices to match others'. This, of course, will force many enterprise shops to stick with their over priced incumbent providers, because a move to EC2 would require answering the question, "why didn't go to IaaS and AWS 5 years ago." Oddly enough, it is easier to hide behind higher prices.

Originally published by VMBlog. For the original post, visit <http://vmblog.com/archive/2013/12/06/cohesiveft-2014-predictions-amazon-will-lead-but-the-gap-is-closing.aspx#>. All rights reserved by the author.

Exclusive Prediction: Cloud Computing in 2015

Patricia B. Seybold

CEO & Sr. Consultant, Patricia Seybold Group



[Excerpt] Customers want to be able to access and manage their own information and activities from anywhere, and they want that information to be securely backed up and redundant. The most cost-effective way to satisfy customers' requirements will be to take advantage of low-cost cloud computing services. To ensure the integrity of customers' data and

communications, large organizations will provide their own application security, data security, and virtual private networking. Smaller firms will rely on VPN and cloud security experts.

Your Customers' Data NEEDS to be in the Cloud

Remember, your firm doesn't own your customers' information. Your firm's *customers* own their own information. They have simply entrusted it to your organization. And, unless you make it easy for them to access, synchronize, and manage their stuff, you're going to lose those customers!

Up until now, enterprise IT architects who have been designing their companies' cloud strategies have decided to leave customer data for last, feeling that, because of its importance and sensitivity, customer data will be the last information to migrate to the cloud. But, if your enterprise architects don't take your customers' needs into account as you plan and execute your firm's migration to private, public, and hybrid clouds, the chances are good that you'll wind up with a "customers last" strategy—one that will keep you from being competitive. If so, there's a good chance that one or more of your competitors will leverage cloud computing and mobile apps to offer your customers easy, secure access to all the information those customers need to keep on top of things, to get things done quickly and easily, and to keep everyone on the same page.

If your IT architects ignore the need to plan your "customer cloud" strategy, customer-facing applications will continue to be developed and deployed by different groups within your organization. Some of these applications will run on proprietary home-grown systems; some will be deployed via the Internet; others will be apps downloaded on customers' mobile devices. Your organization will wind up with an unmanageable, brittle, and fragmented set of mobile apps, customer portals, and intelligent networks and systems that do not represent your brand well. Customers will not have an integrated view that lets them manage their

relationships with your firm and keep track of the products they've purchased from you or that you are managing for them. Customers certainly won't have confidence that their data is consistent and up-to-date.

Therefore you should *lead* your cloud strategy with customer data and applications hosted in secure private clouds and accessed by secure private networks. You'll be able to be more responsive to customers' needs to have their own information at their fingertips from their mobile phones and tablets. You'll be able to provide customers with a rich environment for managing and interacting with the assets they've purchased from you. It will be easier to keep customers' mobile devices, your customer portals, and their intelligent products in synch.

Your Customer Cloud Needs to Be Highly Secure

Of course, when dealing with your firm, your customers will also expect and require more privacy and security from you than they do from Facebook or Google. They expect you to keep their personal and transactional data away from hackers who may be intent on stealing their identity or on compromising their data. They count on you to protect their corporate information from industrial espionage. They also expect you to keep both their personal and corporate information free from spying by government entities. When customers are dealing with you, they expect you to protect their privacy and to protect the integrity of their corporate information. To deliver the levels of privacy, security and data integrity your customers require to manage their activities, you'll need highly secure cloud computing environments and highly secure and encrypted networks to enable your customers to access their information in your customer cloud(s).

These security requirements go beyond the need for physical security of the hosting centers used to house all the cloud computing infrastructure. You will also need to go beyond intrusion detection and anti-hacking measures to protect your cloud computing infrastructure and cloud computing platforms. Customers will expect you to secure and protect their data and network communications at the highest levels—at the application layer, across multiple clouds. And, since many customers' activities also cross organizational boundaries, we'll need security in our application-to-application interactions across multiple clouds as well.

What Will CustomerS Require from Your Cloud by 2015?

By 2015, You'll Need a Secure Customer Cloud and Secure Virtual Private Networks to Be Competitive

Because customers will expect and demand you to manage and synchronize "their" information, they will want you to keep it secure and not able to be viewed or accessed by anyone they haven't authorized to do so.

Encrypt and Secure Customers' Data. You'll need to use the same data security practices in the Cloud that you've used in securing customer data inside your firewall. Your cloud computing environment needs to be password protected, SSL-secured, and your customers' data should be encrypted.

Provide Secure Network Access. Whether your customers are accessing your cloud-hosted applications from their mobile devices and/or from their office systems or from a coffee shop, they should have the benefit of secure access to their own data. Your customer cloud strategy needs to include provisions for both identity management and access management as well as to be secure from unwanted monitoring. You don't want someone parked outside your customer's home to be able to eavesdrop on that customer's supposedly secure interactions with your firm.

Manage Your Own Security Perimeter across Clouds. It doesn't matter to customers whether you are using IBM's or HP's cloud services or Amazon's or Google's—or all of them. Customers will expect your company to ensure that a) their information is backed up and replicated in multiple clouds managed by different cloud providers and b) that you have secured the virtual perimeter around their information and the applications they use to interact with their information across all of those clouds.

By 2015, Customers Will Want Control Over the Jurisdiction in Which Your Customer Cloud(s) Reside

The nice thing about digital services in the cloud is that they can be hosted anywhere in the world. But that's also a problem if the country in which the data actually physically resides has laws that restrict what information can be hosted there. For example, copyrights are national. If I buy a digital book with an American copyright, or download an Australian movie, or stream music from Korea, where am I legally allowed to store and access that information?

Each country and region also has different laws about what customer information can be moved across country boundaries. Even if that information is hosted in the cloud, the digital bits are still stored on a computer disk in one or more physical locations. German customers' information needs to stay in Germany. It can't cross borders, so you still need to be able to

prove that the physical data storage is in Germany. What if my German customer information winds up on a server in India? That's a big problem!

Ensure that Customers' Regulated Data Does Not Cross Borders. Customers will be relying on your company to abide by national and international data privacy regulations. You need to be able to demonstrate to customers that their data is physically resident only in the jurisdictions in which they choose for it to reside.

Government security policies also have cross-border implications. The U.S. Patriot Act apparently gives the US government the ability to access private information from any American company anywhere in the world if terrorist activity is suspected. So, if you are a U.S.-based company and you have German or Brazilian customers, their information may be subject to search by the US government. In order to be competitive in a global market, US companies need to be able to assure their non-US customers that their information is not subject to search by US authorities simply because the company with which they do business is headquartered in the U.S.

By 2015, Customers Will Want to Manage Their Activities in Your/Their Customer Clouds

Customers may not explicitly ask you to provide them with "customer clouds." But if you take an inventory of what your customers actually want and need to do in and around "their" stuff, you quickly realize that you'll need a customer cloud strategy and implementation in order to satisfy all of these customer requirements:

Make It Easy for Customers to Manage Their Stuff (in the Cloud). Today's customers expect to be able to see and manage their relationships and preferences. They also expect to have access to relevant information about everything they've bought either directly from you and/or from one of your channel partners. If they need to re-order, replenish, repair, get help with, buy supplies, and/or get better performance or learn how to do something new, they expect you to provide them with a set of tools that let them get things done. Customers want to keep track of their assets and their projects. They want to be able to do this from anywhere.

Keep Customers' Stuff Up-to-Date, Backed Up, and In Synch. Although we want to have everything at our fingertips, most of us do not have the time and inclination to do data entry. As customers, we don't want to have to enter everything. We want you, the company from whom we bought our assets, and/or the service provider or custodian to whom we've entrusted the care of our assets, to "automagically" pre-populate our stuff and to keep our stuff up-to-date for us as things change. (Show me my current bank balance and the value of

my investment portfolio, show me the current configuration of the hardware and software in my networks, show me what courses I've completed and what credits I've earned to-date, show me the current working components in my power plant, show me the current version of this file, and maintain all the previous versions in case I need to roll back). If we, as end-users, have made changes that you can't detect or know about, we're willing to provide those updates (but not necessarily immediately; only when we happen to look at things and realize that something's not quite up-to-date). Best of all possible worlds, we'd like it if you could automatically keep everything up-to-date for us without our having to DO anything. And, if we're entrusting our assets to your care (or even the records about our assets), we want you to provide automatic back up so we don't have to worry about losing them.

You'll Need to Have Implemented Your Customer Cloud Strategy by 2015

The customer requirements we've described are not new. These are the behaviors that customers are already exhibiting as they interact with the mobile devices and with the Web. But, if you don't step back and realize that you'll need a comprehensive customer cloud architecture to actually meet your customers' expectations within the next two years, you'll be so far behind the next generation of competitors who assume cloud computing as the standard deployment model that you'll never be able to catch up.

Exclusive Prediction: Little Cloud Quantum Houses

Patrick Kerpan

CEO, CohesiveFT



Cloud to the ground. The cloud seems like a cloud (as opposed to also fog and flood) because we don't yet have ubiquity of connectivity, storage, and computation close to us. It is "out there", "up there", across our Internet connection.

However, as we get additional orders of magnitude of connectivity, storage and computation and it gets closer to us - what are we going to call it? When I have more addressable storage and computation in my wearable items than I traditionally have had in my home - what will we call it. Good gosh, I hope to never hear the term "body cloud".

It is cloudy precisely because it is far away. And as we reach the end of what might be considered the first generation public cloud, we can see that the design center of "far away" has taken root in most (if not all) of the offerings. "Far away" has taken root in the feature set and also (IMHO) in the mindset of the providers.

Traditional data center connectivity vendors (metal makers) did not believe the emergence of the hypervisor switch, the virtual network controller, etc was even close to happening. We (CohesiveFT team) got laughed out of Cisco, once in summer of 2008, again in 2010, as we were told (to paraphrase) "no one will run networks via virtual infrastructure, this is a hardware thing".

Now look where we are. The metal makers now know virtual network infrastructures: NFV, SDN, controllers, Open Source hardware, etc.. are all here to stay, and are profoundly changing the nature of their business.

Yet there is a new "we will do it all" crowd, and somewhat surprisingly it is the datacenter, virtual infrastructure, hypervisor infrastructure folks.

As context, CohesiveFT's business is hybrid networking which occurs as part of application deployment in virtual infrastructure. We run in vms/containers that are deployed into the compute fabric. Our wires are software, our boxes are software. We are connectivity that is formed dynamically out of compute. We run far above the hypervisor just like the applications do, providing distributed control, federation and security within the application layer, using hypervisor networking as a service provider to us and a source of "bulk transport" for our IP bits. Where the application can go, we can go.

Despite having a portfolio of 500+ great, connected customers, many of them well known, recognizable industry leaders, who have driven about 100 million hours of virtual network use-cases, many of our hypervisor/virtual data center/public cloud partners tell us "thanks for being you, for now, but really we will do it all".

As the underlying "metal" resources have become bulk compute, bulk storage, and bulk transport consumed by the cloud management platform, they are now "the big dog". Just like the metal makers who said "it will never be virtualized", the virtualizers tell us, "it will never be containerized, distributed and out of our control".

Sorry.

Yes.

It.

Will.

So despite the self proclaimed apotheosis of each layer of the computing stack - the future will bring stacks upon stacks. In the comedic style of our universe it looks more like the Cat in the Hat with Little Cat Z, then it does an ordered, controlled architecture.

It is cloudy precisely because today it is so far away.

But fog and flood are on their way to being close, and closer still. Optics, connectivity, computation and storage are coming down from the Mount, with these Promethean gifts, it will be cloud-to-the-ground.

(Did you notice I added optics to the cloud? In a centrally controlled, "up there" architecture it doesn't exist, but in cloud to the ground optics and mobility are part of the fog and the flood).

We MIGHT have real, practicable quantum computing in the next decade.

We MIGHT have aesthetically pleasing, affordable, high performance solar panels in the form of roof tiles in the next decade.

We MIGHT have powerless network connectivity in the next decade.

We MIGHT have spray-on optical surfaces that look like wall paint but can display or capture images.

All this along with the monotonous drum beat of the doubling of conventional compute power, the more than doubling of conventional storage, etc. every two years. We are into the computing asymptote with conventional techniques, much less the effects of the advances above.

This will mean the roof, the walls, the doors, the floors of my house will be for all intents half data center and half organism. So will my car. So will my clothes.

Far, far, far away from the centralized hypervisor. Applications will swarm and will need self organizing connectivity, federation, management and control. This means they bring their network with them.

Will CohesiveFT and its dynamic, application-centric networks flow from the cloud to the ground? Maybe. But will cloud-to-the-ground happen? Definitely. Will the design center of "far away" prevail? It can't. Each layer of the cloud-to-the-ground stack is going to bring its own quirks, breathy passion, and unswerving advocates.

Ubiquity will be the new design center. Ubiquitous digital atmosphere the air we breathe.

Copyright © 2014 by CohesiveFT.
ISBN-13 978-0-615-95611-4